

PB96-148770

NTTS
Information is our business.

CONVERGENCE BOUNDS FOR MARKOV CHAINS AND APPLICATIONS TO SAMPLING

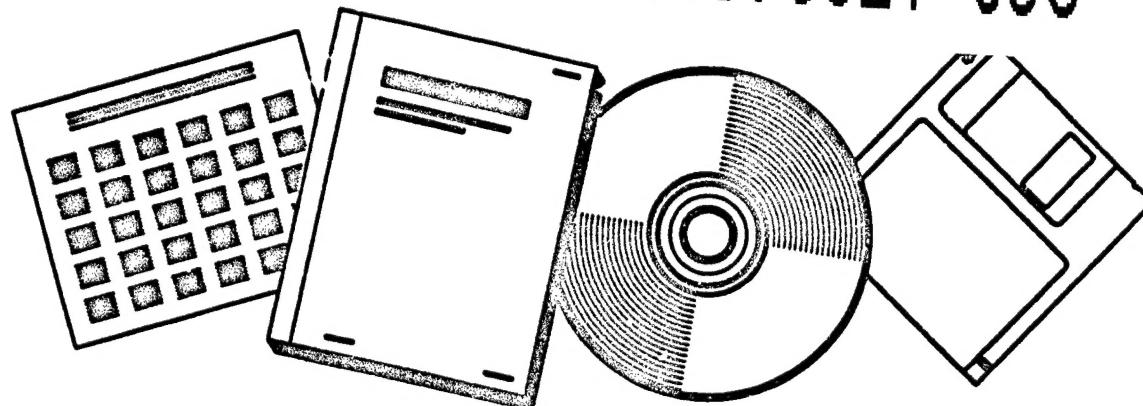
STANFORD UNIV., CA

DISTRIBUTION STATEMENT A

Approved for public release.
Distribution Unlimited

MAY 91

19970821 056



**U.S. DEPARTMENT OF COMMERCE
National Technical Information Service**

DTIC QUALITY INSPECTED 3

May 1991

Report No. STAN-CS-91-1361

thesis



PB96-148770

**Convergence Bounds for Markov Chains and
Applications to Sampling**

by

Anil Ramesh Gangolli

Department of Computer Science

**Stanford University
Stanford, California 94305**

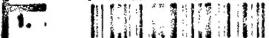


REPRODUCED BY: **NTIS**
U.S. Department of Commerce
National Technical Information Service
Springfield, Virginia 22161

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 6700-0127

PB96-148770



1. REPORT DATE

2. REPORT TYPE AND DATES COVERED

4. TITLE AND SUBTITLE

Convergence Bounds for Markov Chains and Application
to Sampling

6. AUTHORS

Anil Ramesh Gangolli

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Computer Science Department,
Stanford University8. PERFORMING ORGANIZATION
REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

DARPA / ONR
1400 Wilson Blvd.
Arlington VA 2220910. SPONSORING/MONITORING
AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT

UNLIMITED

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

Consider a discrete-time ergodic Markov chain on a finite state space S with stationary distribution π . By simulating such a chain, it is possible to draw random samples from S that have distribution π or nearly π . This thesis treats some basic questions that arise when one wants to apply such a sampling method in a rigorous way.

We begin by reviewing recently developed techniques for proving convergence bounds for Markov chains, and give some new convergence bounds for a number of chains related to "urn models." We then exhibit tight spectral bounds on the variance of natural mean-value estimators computed from a time-reversible Markov chain, and we use these bounds to study issues of efficiency when computing mean-value estimates by this method. Combining the variance bounds with a construction of expander graphs, we obtain an efficient pseudo-random generator for mean-value estimation. Finally, we present some experimental results obtained using the Markov chain sampling method on a statistical problem.

14. SUBJECT TERMS

Analysis of Algorithms, Combinatorial Mathematics, Statistics

15. NUMBER OF PAGES

153

16. PRICE CODE

17. SECURITY CLASSIFICATION
OF REPORT18. SECURITY CLASSIFICATION
OF THIS PAGE19. SECURITY CLASSIFICATION
OF ABSTRACT

20. LIMITATION OF ABSTRACT

CONVERGENCE BOUNDS FOR MARKOV CHAINS AND
APPLICATIONS TO SAMPLING

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF COMPUTER SCIENCE
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

By
Anil Ramesh Gangolli
May 1991

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.

Persi Diaconis
(Principal Advisor)

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.

Donald E. Knuth

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and in quality, as a dissertation for the degree of Doctor of Philosophy.

Rajeev Motwani

Approved for the University Committee on Graduate Studies:

Dean of Graduate Studies

Abstract

Consider a discrete-time ergodic Markov chain on a finite state space S with stationary distribution π . By simulating such a chain, it is possible to draw random samples from S that have distribution π or nearly π . This thesis treats some basic questions that arise when one wants to apply such a sampling method in a rigorous way.

We begin by reviewing recently developed techniques for proving convergence bounds for Markov chains, and give some new convergence bounds for a number of chains related to "urn models." We then exhibit tight spectral bounds on the variance of natural mean-value estimators computed from a time-reversible Markov chain, and we use these bounds to study issues of efficiency when computing mean-value estimates by this method. Combining the variance bounds with a construction of expander graphs, we obtain an efficient pseudo-random generator for mean-value estimation. Finally, we present some experimental results obtained using the Markov chain sampling method on a statistical problem.

Acknowledgements

*When I'm not thank'd at all, I'm thank'd enough,
I've done my duty, and I've done no more.*

—Henry Fielding (1707-1754)

*How sharper than a serpent's tooth it is
To have a thankless child!*
—William Shakespeare (1564-1754), *King Lear*

In 1985, Persi Diaconis gave a truly inspiring series of lectures that introduced to me many of the ideas that underlie this work. He has since been my faithful advisor, friend, hero, and academic father. I cannot thank him enough for his patient guidance.

Don Knuth is the reason I flunked my comprehensive exams on the initial attempt. He led a programming and problem solving seminar that I attended during my first year here. It was so fun, so fascinating, and so engrossing that I spent all my spare time thinking about those problems and neglected to study any of my weak "comp" areas. (With more courses like that, we wouldn't need exams!) Over the period in which I did this work, Don not only provided me with continuing inspiration and little bits of great advice, he arranged my research funding and diligently kept the 'Black Friday' committee at bay. His careful proofreading of this text has prevented several errors from reaching publication.

Rajeev Motwani graciously agreed, when asked somewhat late in the game, to be the third reader on my thesis. I would particularly like to thank him for his suggestions for improving Chapter 4.

David Aldous, Andrei Broder, Pang Chen, Jim Fill, Arif Merchant, and Alistair Sinclair have all shared their ideas on Markov chain methods with me. I really enjoyed the company and support of this hard-working community, especially during our summer seminars at Stanford. I would also like to thank Tomás Feder, Seffi Naor, and Yossi Asar for discussing various ideas with me.

The National Science Foundation (grant numbers DCR-83-51757 and CCR-8610181-A2) and the Office of Naval Research (grant numbers N00014-87-K-0502 and N00014-88-K-0166) sustained various parts of this research. Some of the computing facilities used in this work were donated to

the Computer Science Department by AT&T, by Digital Equipment Corporation, and by IBM.

Because this marks the culmination (I think) of my formal education, it is also appropriate here to thank the numerous other teachers who have guided me along the way. I'd like particularly to thank Jim Erikson, E. Lee Stout, Tom Cover, and Bob Floyd. Ron Graham and Oren Patashnik deserve special mention for exciting me about discrete mathematics with their brilliant performance of CS151 *Concrete Mathematics* in 1981, my first undergraduate year here.

Since $1991 - 1984 = 7$, it takes little imagination to guess that I was not working on my thesis research all of the time. Critical at the other times were my good friends Thane Plambeck, C. Greg Plaxton, Yossi Friedman, Ashok Subramanian, John Woodfill, Andrew Kosoresow, Evan Cohn, Ramsey Haddad, Sherry Listgarten, Gidi Avrahami, R. Michael Young, and the rest of the su.roger-or-andy gang. Gee, what a *smart* bunch! I would especially like to thank two very special friends, Elisabeth Jaffe and Carrie Shook, who gave me immeasurable emotional support and encouragement during these years.

Finally, let me record a necessarily inadequate "thank you" to my parents, Ramesh and Shanta Gangolli, and to the rest of my family for all of their support and faith.

Anil Gangolli
May 1991 at Stanford, California.

Introduction

Everyone understands, intuitively, that when we shuffle a deck of cards, we do so in order to get a random ordering of the cards. We are using a (presumably) stochastic process whose states are orderings of the cards to get to an ordering that is random and does not depend on the original order of the cards.

At least as early as 1953, Metropolis *et al.* [MRR⁺53] suggested the following similar computational technique for sampling from a set. The distribution of an ergodic Markov chain on its state space converges to a unique stationary distribution π that does not depend on the initial distribution. So to sample according to π , simulate a Markov chain with stationary distribution π until it is near stationarity, and then draw samples from the chain.

To use this method in a rigorous way, one must answer questions like: 'How long does it take for the distribution of the chain to reach (or be near) the stationary distribution π ?' 'How do samples drawn from the chain perform (compared to independent π -distributed samples)?'

The classical theory of Markov processes offers little in the way of non-asymptotic answers to these questions. The role of the spectrum of the transition matrix in answering such questions has been known for some time, but until recently this relation has not been very useful because of the lack of good bounds on the spectrum. It is only in the last ten years that people have developed useful techniques to answer these questions for the large sparse chains that one faces in practice. This thesis presents some contributions to this field.

In the first chapter, we present an overview of some new techniques for bounding convergence rates of Markov chains, and in particular for the class of time-reversible Markov chains. This includes the class of random walks on undirected graphs, which have numerous applications in theoretical computer science. We devote particular attention to spectral bounds, based on recent geometric arguments to bound the second largest eigenvalue. These bounds relate connectivity and 'expansion' properties of the underlying graph to bounds on the eigenvalues, and therefore to bounds on the convergence rate.

In the second chapter we use these techniques to prove bounds for two classical 'urn models,' and some direct generalisations. These lead us naturally to consider random walks induced by group actions. There has been a good deal of success recently in using harmonic analysis to obtain

convergence bounds for such processes. In order to be tractable, that type of analysis requires special properties beyond that of the group or group action. At the cost of obtaining poorer bounds, we take a different approach that does not require special additional structure. We obtain new diameter-based bounds on the expansion of the Cayley graph of a group action. This yields diameter and volume based bounds on the convergence rate of certain random walks based on groups.

Estimating the mean value of a function on a set is arguably the most common application of sampling. In Chapter 3 we investigate the natural mean-value estimators on the Markov chain: sample means based on samples drawn some t steps apart from the stationary (or near-stationary) chain. We prove tight worst-case bounds on the variance of such sample means. It is not hard to see that if t is large, so that we draw our samples as far apart as the time required to reach stationarity from any initial position, then the samples are essentially independent, and have properties approximating independent samples. Drawing samples so far apart, however, seems to 'waste the information' in the intervening states, which we 'pay for' by simulating numerous steps of the chain. However, samples drawn at frequent intervals will be correlated, and will tend to increase the variance of the estimates; which method will be better? The results show something surprising. Independent samples give no smaller variance than a comparable number of samples drawn at a spacing t that is typically *much smaller* than the time required to reach stationarity. The two subsequent chapters are applications of this idea.

In Chapter 4, the results on estimation are applied together with a construction of a family of expander graphs to show that, from a very small number of random bits, one can generate a large set of random binary words that are essentially as good as true independent uniform random words for estimating mean values of real-valued functions.

In our fifth chapter, we consider a problem that arises in the analysis of multivariate statistical data: estimating the significance of two-way 'contingency tables' under the uniform distribution. Contingency tables are arrays with nonnegative integer entries whose row and column sums are prescribed values. These tables arise in a surprising variety of settings in the theory of the symmetric group and in statistics. They have drawn considerable attention from combinatorialists and statisticians. It is unknown how to count the exact number of tables with given row and column sums, which suggests that one will not easily find a traditional method of sampling from the set.

We suggest a random walk that provably converges to the uniform distribution on the set. This yields algorithms to estimate significance values and to approximately count the set. The algorithms will give good performance in polynomial time provided that the eigenvalues of the chain can be suitably bounded.

We are able to prove polynomial-time bounds for a slight modification of our suggested algorithm on a certain well-behaved class of instances. But these bounds do not hold for all instances, and are too weak to imply practical running times.

We conjecture that better bounds actually hold, based on computational experience, and also

provide some theoretical motivation for these conjectures. The accuracy of various significance estimates obtained using the walk compare well to values obtained by exact computations for some cases, and also to results obtained by other means of approximation. These experiments demonstrate the practicality of the suggested techniques and show that the walk seems, in practice, to display the conjectured convergence properties.

Notation and Conventions

We use the following notations and conventions frequently. Note, in particular, the definition of graphs that we use.

Notation

\mathbb{R} denotes the real numbers.

\mathbb{Z} denotes the integers, \mathbb{Z}_m the integers modulo m , and \mathbb{N} denotes the nonnegative integers.

$|z|$ denotes cardinality if z is a set, absolute value if z is a number, length if z is a string.

$[n]$ denotes the set of positive integers $\{1, 2, 3, \dots, n\}$.

$\ln n$ denotes the natural (base e) logarithm, while $\lg n$ denotes the logarithm base 2.

U usually denotes the uniform distribution on the set under discussion, usually the state space V of a Markov chain. $U(v) = 1/|V|$ for each $v \in V$.

$P(A)$ when P is a probability distribution on S , and $A \subseteq S$, denotes the total probability of the subset A : $P(A) = \sum_{a \in A} P(a)$.

$\|P - Q\|$ is the total variation between the distributions P and Q . Total variation is a metric on the space of probability distributions on a set. See Appendix A.

$\text{rpd}(P, Q)$ denotes the relative pointwise distance of P from Q . See Appendix A

$\text{sep}(P, Q)$ denotes the separation distance of P from Q , another measure of distance between distributions. See Appendix A.

$P_{ij}^{(k)}$ denotes the (i, j) entry of the matrix P^k . That is, first raise P to the power k , then take the (i, j) entry. The notation P_{ij}^k without the parentheses means $(P_{ij})^k$, the (i, j) entry raised alone to the k th power.

π_k the k th-step distribution of the Markov chain under discussion, given by $\pi_k = \pi_0 P^k$.

$L(P)$ the Laplacian operator associated to the chain P , namely $L = I - P$. We simply write L when P is understood from context.

$Q(G)$ for a graph G (see below) is the matrix $Q = D - A$, where D is the diagonal matrix of degrees $D_{vv} = \deg v$, $D_{vw} = 0$ for $v \neq w$, and A is the adjacency matrix of G . This is the graphical Laplacian that is used by Alon [Al086] and others in the purely graph-theoretic setting. We will simply write Q when G is understood from context. When G is a d -regular graph, and P the natural random walk on G , the Laplacian $L(P)$ is related to $Q(G)$ via $L(P) = \frac{1}{d}Q(G)$.

$(\phi L, \psi)_\pi$, for real-valued functions on V , and measure π with support everywhere on V , denotes the quadratic form based on L under the inner product $(\phi, \psi)_\pi = \sum_v \phi(v)\psi(v)\pi(v)$. That is, $(\phi L, \psi)_\pi = \sum_v [\phi L](v)\psi(v)\pi(v)$. If L is the Laplacian of a reversible ergodic chain, then L is self-adjoint in this inner-product space.

λ_1 denotes the second-largest eigenvalue of the transition matrix for the chain under discussion.

λ_* is the second largest among the *absolute values* of the eigenvalues of the chain under discussion. This should not be confused with $|\lambda_1|$. The two sometimes, but not always, coincide.

μ_1 denotes the smallest strictly positive eigenvalue of the Laplacian $L(P)$ for the chain under discussion

ν_1 denotes the smallest strictly positive eigenvalue of $Q(G)$ for the graph under discussion, generally the underlying graph of a reversible chain.

$\text{Tr}[M]$ for a matrix M denotes its trace, which is the sum of its diagonal entries, and this is equal to the sum of its eigenvalues.

Σ_{rc} is the set of all $m \times n$ nonnegative integer tables with row sums $r = (r_1, r_2, \dots, r_m)$ and column sums $c = (c_1, c_2, \dots, c_n)$, where $\sum_i r_i = \sum_j c_j = N$.

$\text{Binomial}(n, p)$ denotes the binomial distribution with parameters n and p . The discrete random variable X is distributed $\text{Binomial}(n, p)$ when $\Pr\{X = k\} = \binom{n}{k} p^k (1-p)^{n-k}$. For further background, see [Fel70, Vol. 1, Chapter VI].

Other Conventions

We use boldface to mark the definition of a new term. When there is little danger of confusion, we omit commas between multiple subscripts. Thus P_{xy} denotes $P_{x,y}$.

Lemmas, theorems, examples, and figures, are numbered in a common sequence within chapters. So Example 2.1 would precede Lemma 2.2, which would precede Theorem 2.4, which would precede

Figure 2.5. All would be found in Chapter 2. Numbered equations are numbered in a separate sequence within chapters.

We use a slightly nonstandard definition of a graph. For us a graph means a finite undirected multigraph, with self loops allowed. We call a graph simple if it is a graph in the usual sense, without self-loops or multiple edges. Every graph G is naturally associated with a nonnegative integer symmetric matrix A , where A_{ij} is the number of edges between the vertices i and j . This is the adjacency matrix of G . For a vertex v , its degree is defined $\deg v = \sum_j A_{ij} = \sum_i A_{ij}$. Note that this counts self-loop edges only once. This means that $|E| \leq \sum_v \deg v \leq 2|E|$, where we consider E as the multiset of edges. The latter is an equality precisely when the graph has no self-loops. We call a graph regular if every vertex has the same degree. For d -regular graphs without self loops we have $d|V| = 2|E|$. We call a graph bipartite if it has no odd-length cycles. A self-loop constitutes an odd-length cycle by itself.

In a graph $G = (V, E)$, if $A \subseteq V$, then

$$\text{Nbd}(A) = \{w \mid w \notin A \text{ and } (\exists v \in A)[(v, w) \in E]\}.$$

This is the set of vertices that are neighbors of vertices in A but are not themselves in A . Note that a vertex appears only once in $\text{Nbd}(A)$, though multiple edges from A may reach it.

Contents

Abstract	iv
Acknowledgements	v
Introduction	vii
Notation and Conventions	x
1 Techniques for Bounding Convergence Rates	1
1.1 Ergodic Markov Chains	1
1.2 Time-Reversibility	3
1.3 Random Walks on Graphs	5
1.4 Convergence in Terms of the Spectrum	7
1.4.1 Bounds for Reversible Chains	7
1.4.2 Strong Aperiodicity	9
1.4.3 Reversibilization	9
1.5 Geometric Eigenvalue Bounds	11
1.5.1 The Laplacian	11
1.5.2 Cheeger-Type Bounds	14
1.5.3 Canonical Path Arguments	16
1.5.4 Poincaré-type Bounds	19
1.6 Probabilistic Bounds	22
1.6.1 Couplings	22
1.6.2 Strong Stationary Times	24
2 Urn Models and Group Actions	28
2.1 Ehrenfest-type Models	28
2.1.1 Arbitrary Steps	28
2.1.2 Adjacent Moves	30

2.2	Bernoulli-Laplace-type Models	34
2.2.1	A Coupling for Bernoulli-Laplace Processes	34
2.2.2	Tight Analysis of the Two-Urn Case	36
2.2.3	Weak Analysis of the General Case	36
2.3	Markov Chains based on Groups	38
2.3.1	Transitive Group Actions	38
2.3.2	Cayley Graphs	40
2.3.3	Vertex-Transitive Graphs	41
2.3.4	Chains Based on Groups	42
2.3.5	On the Harmonic Analysis Approach	45
2.3.6	Magnification Bounds	46
2.3.7	Eigenvalue Bounds for the Chains	49
2.3.8	Examples	50
3	Mean-Value Estimation	53
3.1	Variance Bounds	56
3.2	Sampling to Achieve Given Variance	61
3.3	Indicator Functions	65
3.4	Central Limit Theorem	67
4	Using Expanders in Estimation	69
4.1	Preliminaries	70
4.2	Outline of the Algorithm	72
4.3	Sample Means from G_n	74
4.4	Majorities from G_n	75
4.5	Combining the Results	78
4.6	The Implied Algorithm	78
4.7	Discussion	78
5	Estimating the Significance of Contingency Tables	82
5.1	A Random Walk on Σ_{rc}	85
5.2	Bigger Steps	87
5.3	Eigenvalue Hypothesis	89
5.4	Experimental Results	90
5.4.1	Background	91
5.4.2	Pinckney Gag Rule	92
5.4.3	Hair Color v. Eye Color	94
5.4.4	Irregular Margins	96

5.4.5 Scaling	101
5.5 Provably Polynomial-Time Methods	101
5.6 More on the Eigenvalue Hypothesis	105
6 Directions for Future Work	107
A Notions of Approximation	109
A.1 Point Approximations	109
A.2 Approximate Distributions	110
A.2.1 Total Variation	110
A.2.2 Separation and Relative Pointwise Distance	110
A.2.3 Approximation within Ratio	111
A.2.4 Kolmogorov-Smirnov Distance	112
B Sampling from Near-Stationary Chains	113
B.1 Nearly Independent, Near-Stationary Samples	114
B.2 Other Near-Stationary Samples	115
B.3 On Near-Independence and the Median Lemma	116
C Enumerating Contingency Tables	118
C.1 Classical Counting Approaches	118
C.1.1 Exhaustive Enumeration	118
C.1.2 A Recursive Formula	119
C.1.3 Approximation Formulas	119
C.1.4 Tables and Group Theory	120
C.1.5 Counting with Generating Functions	123
C.2 Hardness of a Related Problem	125
C.3 Approximate Counting using Sampling	127
Bibliography	131

Chapter 1

Techniques for Bounding Convergence Rates

In this chapter we summarise techniques for bounding the rate of convergence of Markov chains. These techniques are applied to get new results in later chapters. This chapter is largely expository, and is not intended to present new contributions of the author, although it contains a few new arguments in some examples.

We first give a brief summary of the basic ergodic theory of Markov chains, and in particular of random walks on graphs viewed as time-reversible Markov chains. Readers that are unfamiliar with this basic theory may find Karlin's text [Kar68] helpful. His appendix covering the Perron-Frobenius theory of stochastic matrices is essential background.

Then we show how upper bounds on the time to convergence can be obtained from bounds on the eigenvalues of the Markov chain. The necessary bounds on the eigenvalues are obtained by geometric means, involving "expansion" and related connectivity properties of the underlying graph of the Markov chain.

Finally we discuss some probabilistic techniques. These are based on the construction of certain stopping time random variables having the property that bounds on the tails of their distributions provide bounds on the variation distance.

1.1 Ergodic Markov Chains

Let V be a finite set of states and let $\{X_k \mid k \geq 0\}$ be a sequence of random variables taking values in V such that for each k the following property holds

$$\Pr\{X_k = v \mid X_0 = v_0, X_1 = v_1, \dots, X_{k-1} = v_{k-1}\} = \Pr\{X_k = v \mid X_{k-1} = v_{k-1}\}. \quad (1.1)$$

A probability distribution π on V is called a stationary distribution of P if it is a left unit eigenvector of the matrix i.e., if $\pi P = \pi$. We can think of such a π as a fixed point of P , and we might then expect convergence of π_k to such a stationary π under certain conditions. In fact, the following theorem confirms this intuition.

Theorem 1.1 (Basic Convergence Theorem) *Let P be an ergodic Markov chain on V . The following two conditions are equivalent for any probability distribution π on V :*

1. π is a stationary distribution of P : $\pi P = \pi$;
2. for every $v \in V$, $\lim_{n \rightarrow \infty} \pi_k(v) = \pi(v)$ regardless of the choice of π_0 .

Furthermore, there exists a unique distribution π satisfying these conditions. This distribution is nonzero everywhere on V .

This classical theorem tells us that for an ergodic chain, no matter how we choose the initial distribution, eventually the distribution of the state approaches a unique stationary distribution on the state space. For brevity, we say simply that the process, rather than 'the distribution of the state of the process,' converges to its stationary distribution.

A Markov chain P is called doubly stochastic if P^T , the transpose of the matrix P , is also a stochastic matrix. This is the same as saying that not only the rows, but also each of the columns sums to 1. A chain P is called symmetric if $P = P^T$. Note that a symmetric chain is necessarily doubly stochastic. The uniqueness of the stationary distribution gives the following corollary.

Corollary 1.2 *If P is an ergodic chain on S , it converges to the uniform distribution on S if and only if P is doubly stochastic. In particular, if P is symmetric, then it converges to the uniform distribution on S .*

The theorems of this section do not tell us anything about how fast the convergence will be. We would like bounds on how far the k th-step distribution π_k will be from the stationary distribution π . We spend the rest of this chapter describing techniques to address this problem.

1.2 Time-Reversibility

Recall that the k th-step distribution of a Markov chain with transition matrix P is given by $\pi_k = \pi_{k-1}P = \pi_0P^k$, where π_0 is the initial distribution. In general, the action of a transition matrix P on the state distribution π_k is hard to analyse. However, when the linear operator P can be converted to a matrix diagonalisable over an orthonormal basis of eigenvectors, the action of P is reasonably simple when viewed in this basis. In this section, we develop this idea.

If P is an ergodic chain with stationary distribution π , the time reversal P^R of P is the chain whose matrix entries are $P_{sy}^R = \frac{\pi(y)}{\pi(x)}P_{yx}$, where π is the stationary distribution of P . Intuitively, if

the chain P is supposed to be evolving in the stationary distribution π , then P_{xy}^R gives the probability that P was in state y at time $k-1$ given that it is in state x at time k . Thus when P is evolving in the stationary distribution, P^R is the transition matrix for the process P observed with time reversed,

An ergodic Markov chain P with stationary distribution π is called time-reversible if $P = P^R$. Equivalently, this means that for each pair of states x and y ,

$$\pi(x)P_{xy} = \pi(y)P_{yx}.$$

Many chains that arise from physical problems are time-reversible. In the next section we show that the random walk on any undirected graph is time-reversible, which gives rise to a number of combinatorially interesting chains.

Note that the condition that P is both time-reversible and doubly stochastic (i.e., has uniform stationary distribution) is equivalent to the condition that P is symmetric. But when P is not doubly stochastic, time-reversibility gives a more general symmetry condition.

The algebraic significance of time-reversibility is that it allows us to transform P into a diagonalizable form. Suppose P is an ergodic time-reversible Markov chain on a state space V of n elements with stationary distribution π , and let R be the diagonal matrix with z th diagonal entry $\sqrt{\pi(z)}$. Let $M = RPR^{-1}$. Then

$$M_{xy} = \sqrt{\frac{\pi(x)}{\pi(y)}} P_{xy} = M_{yx},$$

so M is symmetric. Thus we can write

$$M = \Gamma B \Gamma^T,$$

where Γ is the orthogonal matrix whose columns are the eigenvectors of M , and B is the diagonal matrix of the eigenvalues of M , all of which are real; note that these are also the eigenvalues of P .

The Perron-Frobenius theorem insures that exactly one of these eigenvalues $\lambda_0 = 1$, and that all of the other eigenvalues have absolute value smaller than 1. We use $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_{n-1} > -1$ to denote the remaining eigenvalues in decreasing order. We use λ_+ to denote the largest absolute value of any of these non-unit eigenvalues:

$$\lambda_+ = \max\{|\lambda_i| \mid 1 \leq i \leq n-1\} = \max\{|\lambda_1|, |\lambda_{n-1}|\}.$$

We call λ_+ the second absolute eigenvalue of P . This should not be confused with $|\lambda_1|$. They sometimes, but not always, coincide.

The columns of Γ give an orthonormal basis in which we can more easily treat P . Let Γ_z denote the column of Γ corresponding to the state z . Then B_{zz} is the corresponding eigenvalue of P . The values λ_i , $0 \leq i \leq n-1$ are some permutation of the values B_{zz} , $z \in V$. Let $z \in V$ be the state such that $B_{zz} = \lambda_0 = 1$. Note that the eigenvector Γ_z of M corresponding to λ_0 has entries $\Gamma_{yz} = \sqrt{\pi(y)}$, where π is the stationary distribution of P .

We can now state the main reason for working with time-reversible Markov chains.

Theorem 1.3 (Spectral Representation) *Let P be a time-reversible ergodic Markov chain on S with stationary distribution π . Then for all $x, y \in S$, we have*

$$P_{xy}^{(k)} = \pi(y) + \sqrt{\frac{\pi(y)}{\pi(x)}} \sum_{w \neq x \in V} B_{ww}^k \Gamma_{xw} \Gamma_{yw}.$$

Proof: Since $M = \Gamma B \Gamma^T$ and $\Gamma^{-1} = \Gamma^T$ by orthogonality, we have $M^k = \Gamma B^k \Gamma^T$; hence

$$M_{xy}^{(k)} = \sum_{w \in V} B_{ww}^k \Gamma_{xw} \Gamma_{yw}.$$

We also know that

$$P_{xy}^{(k)} = \sqrt{\frac{\pi(y)}{\pi(x)}} M_{xy}^{(k)}.$$

Combine the two. Since $B_{zz} = \lambda_0 = 1$, and $\Gamma_{zz} = \sqrt{\pi z}$ for all z , we find that the term corresponding to $w = z$ gives exactly $\pi(y)$. The other terms appear in the sum. \blacksquare

1.3 Random Walks on Graphs

Most of the results in this thesis deal directly with and apply generally to time-reversible Markov chains. However, many of the chains that arise as examples can be viewed as random walks on graphs. Here we describe the connection.

Let $G = (V, E)$ be a graph. (Our definition is nonstandard; see *Notations and Conventions*, p. xff). The natural random walk on G is the Markov chain with state space V determined by the following process. If the current state is a given vertex v , an edge $\{v, w\}$ is chosen uniformly at random amongst the $\deg v$ possibilities. This chosen edge is then traversed, and the next state is the neighbor w along that edge. The transition matrix for this process is given by:

$$P_{vw} = \begin{cases} \frac{1}{\deg v} & \text{if } (v, w) \in E \\ 0 & \text{otherwise.} \end{cases}$$

Theorem 1.4 *If $G = (V, E)$ is a connected graph that is not bipartite, the natural random walk process P on G converges to the stationary distribution*

$$\pi(v) = \frac{\deg v}{\sum_v \deg v}.$$

Proof: Because G is connected, there is a path of possible transitions between every two states, thus P is irreducible. Provided G is not bipartite, there exist both odd-length and even-length sequences of possible transitions from any state x back to x . This insures P is aperiodic. Since P is both

irreducible and aperiodic it is ergodic. One can easily check that the function $\pi(v) = \deg v / \sum_v \deg v$ is indeed a stationary distribution of the walk.

$$\begin{aligned}
 [\pi P](w) &= \sum_{v:(v,w) \in E} \frac{1}{\deg v} \pi(v) \\
 &= \frac{1}{\sum_v \deg v} |\{v \mid (v, w) \in E\}| \\
 &= \frac{\deg w}{\sum_v \deg v} \\
 &= \frac{\deg w}{\sum_w \deg w} \\
 &= \pi(w).
 \end{aligned}$$

By Theorem 1.1, we know that this is the unique stationary distribution and that the distribution of the position of the walk will tend to this distribution. ■

Note what this says informally: in the long term, the proportion of time that the walk spends on vertex v is proportional to the degree of the vertex.

Corollary 1.5 Suppose G is any connected graph that is not bipartite. Then the stationary distribution of the random walk on G is the uniform distribution on V if and only if G is regular.

Theorem 1.6 If G is a connected non-bipartite graph, the natural random walk on G is ergodic and time-reversible.

Proof: Let P denote the walk process on $G = (V, E)$. By Theorem 1.4, the stationary distribution of P on V is $\pi(v) = \deg v / \sum_v \deg v$, and by definition, $P_{vw} = \frac{1}{\deg v}$. Now note, that since (v, w) is an edge only if (w, v) is:

$$\pi(v)P_{vw} = \frac{1}{\sum_v \deg v} = \frac{1}{\sum_w \deg w} = \pi(w)P_{ww}.$$

■

There is a much weaker, but useful, relation in the converse. Suppose P is ergodic and time-reversible on V . There is an associated undirected graph on V such that every transition of P corresponds to a step along an edge of the graph. Suppose \bar{E} denotes the set of possible directed transitions $e = (x, y)$ with $P_{xy} > 0$. These come in pairs, since $(x, y) \in \bar{E}$ iff $(y, x) \in \bar{E}$ follows from reversibility and ergodicity. Let E denote the set of possible transitions, $\{x, y\}$, $x \neq y$, taken without orientation. We can define an undirected simple graph $G_P = (V, E)$, sometimes called the underlying graph of the chain. This graph is simple; it has no self-loops or multiple edges. If P is the random walk on a graph G , then G_P is obtained from G by removing self-loops and collapsing multiple edges to single ones.

This underlying graph will play a large role in the geometric eigenvalue bounds presented later in this chapter.

1.4 Convergence in Terms of the Spectrum

In this section we discuss bounds on the convergence rate based on the eigenvalues of the transition matrix. Later we discuss methods of bounding these eigenvalues.

In order to speak precisely about the convergence rate, we first need to introduce distance measures between distributions. We will generally use one of the following two notions of distance between distributions on V . The total variation, $\|\pi_k - \pi\|$ is defined

$$\|\pi_k - \pi\| = \sup_{A \subseteq V} |\pi_k(A) - \pi(A)|. \quad (1.2)$$

It is commonly used in statistical settings. Another notion of distance is handy for its nice combinatorial properties. For $r > 0$, we say that the distribution π_k approximates π within ratio r if for all points $v \in V$, $\frac{1}{r}\pi_k(v) \leq \pi(v) \leq r\pi_k(v)$. The latter is closely related to the relative pointwise distance

$$\text{rpd}(\pi_k, \pi) = \max_{v \in V} \frac{|\pi_k(v) - \pi(v)|}{\pi(v)} \quad (1.3)$$

used by Jerrum and Sinclair in their work on this topic [SJ87] [JS88] [JS90]. Further important background material appears in Appendix A.

1.4.1 Bounds for Reversible Chains

We use a lemma to prove our main spectral convergence bound.

Lemma 1.7 (Adapted from [SJ87].) *Let P be a time-reversible ergodic Markov chain with stationary distribution π and second absolute eigenvalue λ_2 . For any initial distribution π_0 on V we have*

$$\text{rpd}(\pi_k, \pi) \leq \frac{\lambda_2^k}{\pi_{\min}},$$

where $\pi_{\min} = \min_{v \in V} \pi(v)$. **Proof:** By the Spectral Representation Theorem 1.3, we have

$$\pi_k(y) = \sum_z \pi_0(z) P_{zy}^k = \pi(y) + \sum_z \pi_0(z) \sqrt{\frac{\pi(y)}{\pi(z)}} \sum_{w \neq z \in V} B_{zw}^k \Gamma_{zw} \Gamma_{yw}.$$

So that

$$\begin{aligned} \frac{|\pi_k(y) - \pi(y)|}{\pi(y)} &= \frac{|\sum_z \pi_0(z) \sum_{w \neq z \in V} B_{zw}^k \Gamma_{zw} \Gamma_{yw}|}{\sqrt{\pi(z)\pi(y)}} \\ &\leq \frac{\max_z |\sum_{w \neq z \in V} B_{zw}^k \Gamma_{zw} \Gamma_{yw}|}{\pi_{\min}} \\ &\leq \frac{\lambda_2^k}{\pi_{\min}}. \end{aligned}$$

Here the first inequality is obtained by extracting the maximum term and using the fact that $\sum_z \pi_0(z) = 1$ in the numerator. The denominator is bounded trivially. Then the second inequality is

obtained by bounding $|B_{ww}^k|$ by $|\lambda_*^k|$, and using the fact that $|\sum_w \Gamma_{zw} \Gamma_{yw}| \leq 1$, by the orthogonality of Γ . \blacksquare

The following theorem puts these in a convenient form, summarizing the λ_* -based convergence bounds.

For the chain P , with stationary distribution π and second absolute eigenvalue λ_* , define the function

$$T_P(\epsilon) = \frac{1}{1 - \lambda_*} \left(\ln \frac{2}{\epsilon \pi_{\min}} \right) \quad (1.4)$$

Theorem 1.8 For any initial distribution π_0 , the following holds for any positive $\epsilon \leq 1$. Whenever

$$k \geq T_P(\epsilon)$$

we have

$$\text{rpd}(\pi_k, \pi) \leq \frac{\epsilon}{2},$$

and

$$\|\pi_k - \pi\| \leq \frac{\epsilon}{4},$$

and

π_k approximates π within ratio $1 + \epsilon$.

Proof: The relative pointwise distance bound here follows from the preceding lemma, by expressing the bound there in terms of the exponential, and using the Taylor expansion for the logarithm. The relative pointwise distance bound implies the variation distance bound by Proposition A.2 in Appendix A. The bound for approximation within ratio follows from the bound on relative pointwise distance and the fact that approximation to within relative error $\epsilon/2$ implies approximation within ratio ϵ . (See Appendix A.) \blacksquare

In general we will call a function $f(\epsilon)$ a convergence guarantee for the chain P if for every initial distribution π_0 , taking $k > f(\epsilon)$ insures $\text{rpd}(\pi_k, \pi) \leq \epsilon/2$, (which implies also that the bounds on variation distance and approximation within ratio given in the preceding theorem hold). Thus, Theorem 1.8 states that $T_P(\epsilon)$ is a convergence guarantee for P , but there may be better guarantees available.

In some cases we may be able to determine the full spectrum of P . When this happens, and the chain P is not only time-reversible, but symmetric, we can get the following bound utilising the full spectrum of P . This version often gives provably tight answers.

Theorem 1.9 (Symmetric Case, Full Spectrum Bound) [AD86] If P is symmetric, then for every initial distribution π_0 we have

$$\|\pi_k - U\|^2 \leq \frac{1}{4} (\text{Tr}[P^{2k}] - 1).$$

These will be our basic tools for proving upper bounds on convergence rates. However, these theorems alone provide no clue of how to bound the spectrum, a problem that in general is hard for large state spaces. We will see techniques for bounding the spectrum in the next section.

1.4.2 Strong Aperiodicity

The necessity of bounding the second largest eigenvalue in *absolute value* poses a minor technical problem, because most of the techniques available for bounding the spectrum only give us information about λ_1 , rather than λ_2 . The following theorems provide a simple method of converting a Markov chain to a time-reversible one with nonnegative eigenvalues, and the same stationary distribution.

A Markov chain P is called **strongly aperiodic** or **diagonally dominant** if $P_{zz} \geq \frac{1}{2}$ for each z , that is, if at each state the probability of immediate return is at least 1/2.

Theorem 1.10 *If P is irreducible and strongly aperiodic then all of its eigenvalues are nonnegative, so that then, in particular, $\lambda_1 = \lambda_2$.*

Proof: Consider the matrix $M = 2P - I$. This is a stochastic matrix that is irreducible, (not necessarily aperiodic). Applying the Perron-Frobenius theorem, the eigenvalues ν_i of M satisfy $-1 \leq \nu_i \leq 1$. The eigenvalues λ_i of P are related to the eigenvalues ν_i of M by $\lambda_i = \frac{1}{2}(\nu_i + 1) \geq 0$.

■

One can get a strongly aperiodic process from any Markov process P as follows. The proof is obvious.

Theorem 1.11 *If P is an irreducible chain then $P' = \frac{1}{2}(I + P)$ is ergodic and strongly aperiodic, and has the same stationary distribution as P . If P is time-reversible, then so is P' . If P is doubly stochastic, then so is P' . If P is symmetric, then so is P' .*

For these reasons, we will sometimes speak of the **strongly aperiodic form** of a chain P , by which we mean the chain $P' = \frac{1}{2}(I + P)$. One can imagine this chain as the following process. At each step flip a fair coin. If the coin comes up 'heads' make a move according to P , otherwise remain in the current state.

1.4.3 Reversibilization

Similarly, we can also make a reversible chain from any ergodic one.

Lemma 1.12 *If P is an ergodic chain, with stationary distribution π , then the stationary distribution of P^R is also π .*

Proof: Simply verify that $\pi P^R = \pi$. For each y , we have

$$\begin{aligned} [\pi P^R](y) &= \sum_x \pi(x)(P^R)_{xy} \\ &= \sum_x \pi(x) \left[\frac{\pi(y)}{\pi(x)} P_{yx} \right] \\ &= \pi(y) \sum_x P_{yx} \\ &= \pi(y) \end{aligned}$$

where the last equality holds because P is stochastic. \blacksquare

Theorem 1.13 *Let P be an ergodic chain with stationary distribution π . Let P^R be the time-reversal of P , and define $P' = \frac{1}{2}(P + P^R)$. P' is ergodic, time-reversible, and also has the stationary distribution π . Furthermore, if P is strongly aperiodic, then so is P' .*

Proof: It is clear that since P is ergodic, so is P' . From the lemma above, π is the stationary distribution of P^R , thus also of P' . Now to show that P' is time-reversible, we need only show that $(P')^R = P'$. It is easy to see that the (x, y) entries are equal:

$$\begin{aligned} [(P')^R]_{xy} &= \frac{\pi(y)}{\pi(x)} P'_{yx} \\ &= \frac{\pi(y)}{\pi(x)} \left[\frac{1}{2} (P_{yx} + (P^R)_{yx}) \right] \\ &= \frac{1}{2} \frac{\pi(y)}{\pi(x)} [P_{yx} + \frac{\pi(x)}{\pi(y)} P_{xy}] \\ &= \frac{1}{2} \frac{\pi(y)}{\pi(x)} P_{yx} + P_{xy} \\ &= (P')_{xy}. \end{aligned}$$

Finally note that $P'_{zz} = P_{zz}$. Thus P' will be strongly aperiodic exactly when P is strongly aperiodic.

\blacksquare

Fill [Fil90] gives some useful theorems relating convergence rate bounds of the non-reversible chain to its reversible version. He also discusses a different multiplicative form of reversibilization, given by PP^R . This multiplicative reversibilization, however, does not always preserve ergodicity.

The following example indicates that reversibilization can lose accuracy, but that there may not be any natural useful spectral convergence bounds in the absence of reversibility.

Example 1.14 The de Bruijn graph of order n is the directed graph with $N = 2^n$ vertices labelled by the integers modulo 2^n , and having a directed edge from the vertex labelled x to that labelled

y if $y = 2x$ or $y = 2x + 1$ each taken modulo 2^n . This can be viewed as left-shifting the binary representation of x , discarding the left-most bit, and placing either a 0 or a 1 in the right-most coordinate. The graph is directed, and the natural random walk on this graph is not time-reversible. It is, however, ergodic. It is not hard to see that one bit gets 'randomized' with each step, so that the position of a random walk on this graph is exactly uniform after exactly n steps. After $n - 1$ steps the distribution is supported on only half of the vertices, so the total variation is large. The n th step yields exactly the uniform distribution. The only non-unit eigenvalues of the transition matrix are all zero, precluding any bound directly involving multiplicative factors like λ_2^k or λ_1^k .

The convergence of the reversible chain $P' = \frac{1}{2}(P + P^R)$, on the other hand, can be related to a symmetric random walk on a line segment of length n , and the time required for convergence of P' can be shown to be $\Theta(n^2)$. Thus, making the chain reversible in this way can cause a significant loss in the rate to stationarity.

Using the multiplicative reversibilization here does not seem to help; the chain PP^R is not even ergodic. \square

1.5 Geometric Eigenvalue Bounds

Some of the better-known ways to get bounds on the eigenvalues of the transition matrix fail in the contexts that interest us. Methods involving direct computation with the matrix (see [GL89]) can give accurate approximations to the eigenvalues for small matrices and when the matrix can be given explicitly, but are not generally useful when we wish to prove bounds for whole classes of chains, where the matrix is only implicitly known or specified, and when the state space is exponentially large in the natural parameters. Likewise, bounds such as those of the Gershgorin type and those described in [CDS80] are typically unusable for the large sparse matrices that we encounter. They tend to give results like $\lambda_* \leq 1$, which we know in any case by the Perron-Frobenius theorem.

In this section we present some bounds that do seem to be useful for the type of problems that we wish to consider. These techniques are essentially geometric. They give bounds in terms of geometric properties of the underlying graph of the chain.

Later, in Chapter 2, we will briefly discuss harmonic analysis, which is a different method that has been used successfully to get bounds on the eigenvalues of certain Markov chains associated with groups.

1.5.1 The Laplacian

Let P be a reversible ergodic Markov chain on a finite state space V with stationary distribution π . We associate to P a Laplacian $L = I - P$, where I is the identity matrix. We will view this as an operator on the functions $\phi : V \rightarrow \mathbb{R}$, which we can also view as (row) vectors in $\mathbb{R}^{|V|}$. (Viewing such a function ϕ as a vector, the action is then given by ϕL .) We impose an inner-product

$(\phi, \psi)_\pi = \sum_v \phi(v)\psi(v)\pi(v)$. For example, if $\mathbf{1} = (1, 1, \dots, 1)$ then $(\phi, \mathbf{1})_\pi$ is the mean-value of ϕ under the distribution π .

Since P is ergodic, we have $\phi L = 0$, the zero function, if and only if $\phi = c\pi$ for some c . The operator L has $|V|$ real eigenvalues μ_i , $0 \leq i < |V|$, related to those of P by $\mu_i = 1 - \lambda_i$. Thus taking the eigenvalues in increasing order and in multiplicity, we have

$$\mu_0 = 0 < \mu_1 \leq \mu_2 \leq \dots \leq \mu_{|V|-1} < 2.$$

Lower bounds on μ_1 can be translated to upper bounds on λ_1 of P via

$$1 - \mu_1 = \lambda_1.$$

In this section, we will discuss bounds on μ_1 . It should be understood that these can then be used directly to bound the convergence rate by applying the theorems of the previous section.

Let \bar{E} denote the set of possible directed transitions $e = (z, y)$ with $P_{zy} > 0$. Recall that the directed transitions come in pairs; $(z, y) \in \bar{E}$ iff $(y, z) \in \bar{E}$ from reversibility and ergodicity. Let E , again, denote the set of possible transitions, $\{z, y\}$, $z \neq y$ taken without orientation. Recall that $G_P = (V, E)$ is called the underlying graph of P and is simple.

Define $F(z, y) = \pi(z)P_{zy} = \pi(y)P_{yz}$. Note $F(z, y)$ is nonzero only when (z, y) is a possible transition. It is a symmetric matrix, by reversibility. For $z \neq y$, we may thus write $F(e)$, for $e \in E$ without ambiguity. It can be thought of as the "flow" of probability mass over the edge e when the chain is evolving in the stationary distribution.

For $e = (z, y) \in \bar{E}$, define

$$\nabla\phi(e) = (\phi(y) - \phi(z)).$$

For $\phi, \psi : V \rightarrow \mathbb{R}$, define the Dirichlet form

$$(\phi L, \psi)_\pi = \sum_v (\phi L)(v)\psi(v)\pi(v) = \frac{1}{2} \sum_{z, y} (\phi(y) - \phi(z))(\psi(y) - \psi(z))F(z, y).$$

This is the quadratic form on L in the space of real-valued functions on V with inner product $(\phi, \psi)_\pi = \sum_v \phi(v)\psi(v)\pi(v)$. One can see that the form is symmetric (L is self-adjoint in this space), and as we already knew, positive semi-definite. Using our notation, we can write

$$(\phi L, \phi)_\pi = \frac{1}{2} \sum_{e \in \bar{E}} (\nabla\phi(e))^2 F(e) = \sum_{e \in \bar{E}} (\nabla\phi(e))^2 F(e). \quad (1.5)$$

In the latter equality, the $\frac{1}{2}$ factor and \bar{E} have been replaced with the set E of undirected transitions. In making this replacement we have used the fact that the undirected edge $\{z, y\}$ is represented twice by directed edges, once by (z, y) and again by (y, z) , each contributing the same value of the term $(\nabla\phi(e))^2 F(e)$, and also that $\nabla\phi(z, y) = 0$ for $z = y$. For the undirected edges in the resulting form, these terms may be computed with either orientation assigned to the edge; the squaring and the symmetry of $F(e)$ makes the choice irrelevant.

Rayleigh's principle gives

$$\mu_1 = \inf \left\{ \frac{(\phi L, \phi)_\pi}{(\phi, \phi)_\pi} \mid (\phi, 1)_\pi = 0, \phi \neq 0 \right\} \quad (1.6)$$

$$= \inf \left\{ \frac{\sum_{e \in E} (\nabla \phi(e))^2 F(e)}{(\phi, \phi)_\pi} \mid (\phi, 1)_\pi = 0, \phi \neq 0 \right\}, \quad (1.7)$$

the infimum being attained precisely at any eigenfunction μ_1 .

This is a weighted form of the Rayleigh quotient of the graphical Laplacian $Q(G_P)$ for the underlying graph G_P . The graphical Laplacian of a graph G is $Q(G) = D - A$ where D is the diagonal matrix of degrees, $D_{vv} = \deg v$, and A is the adjacency matrix of the graph. This is symmetric and positive semi-definite, with smallest eigenvalue 0, and smallest positive eigenvalue ν_1 given by

$$\nu_1 = \inf \left\{ \frac{\sum_{e \in E} (\nabla \phi(e))^2}{(\phi, \phi)} \mid (\phi, 1) = 0, \phi \neq 0 \right\}, \quad (1.8)$$

where (ϕ, ϕ) is the usual L^2 inner product.

This fact can be used to translate known lower bounds for ν_1 to lower bounds for μ_1 . In particular the following property is immediate.

Theorem 1.15 *If P is an ergodic symmetric chain on V then*

$$p\nu_1 \leq \mu_1 \leq p'\nu_1$$

where

$$p = \min_{(x,y) \in E} P_{xy} \quad \text{and} \quad p' = \max_{(x,y) \in E} P_{xy}$$

are respectively the minimum and maximum nonzero probability of any transition (x, y) , $x \neq y$.

Proof: The stationary distribution is uniform on V ; $\pi(v) = \frac{1}{|V|}$. Thus the inner-product $(\phi, \psi)_\pi$ is the usual L^2 inner-product scaled by the multiplicative factor $1/|V|$. Also $F(x, y) = \frac{P_{xy}}{|V|}$ must lie between $\frac{p}{|V|}$ and $\frac{p'}{|V|}$. Plugging these into the quotient of (1.7), bringing out p in the numerator and cancelling $1/|V|$ from both numerator and denominator gives the result. ■

Remark: The relationship of the eigenvalues of $Q(G)$ and connectivity properties of the graph has been known and studied for some time. The eigenvalue ν_1 was called "algebraic connectivity" by Fiedler [Fie73] who showed some relationships to the standard notion of edge-connectivity. The relation between this eigenvalue and expansion properties in graphs has been investigated by Alon and co-authors [Alo86] [AM85] [AGM87] as well as numerous others. These are the Cheeger-type bounds of the next section. The relation of the smallest positive eigenvalue to isoperimetric quantities in the continuous realm was known earlier [Che70]. The Poincaré-type bounds we will see later seem also to have earlier analogues in continuous cases. There seems to be a wealth of deep and interesting

questions to pursue in trying to find other useful connections to known results in the continuous case.

The operators $L(P)$ and $Q(G)$ are discrete analogues of the continuous Laplacian operator. Their Rayleigh quotients in (1.7) and (1.8) are the natural discrete analogues of that arising in the "membrane" problem under "free" (Neumann) boundary conditions, with sums replacing integrals. The Laplacian figures prominently in the diffusion theory of heat, sound, and fluids. Here is some useful intuition: A random walk on a regular graph starts out as a point mass at a vertex of the graph, and the transition matrix P (or alternatively the Laplacian L) serves as the diffusion operator, smoothing the measure to uniformity. This is analogous to the situation of a free membrane that is "poked" at a point, causing a wave-like spreading vibration which eventually settles to uniformity. (This idea can be exploited to give interesting graphic displays of the evolution of random walks on subgraphs of the plane grid.)

The matrix $Q(G)$ also appears in a well-known determinant formula for the number of spanning trees of the graph. $Q(G)$ is also equal to the product CC^T where C is the $|V| \times |E|$ incidence matrix of the vertices to edges of any directed orientation of G . (See [Big74, page 35] and [Knu68, Exercises 2.3.4.2.18-20]) \square

1.5.2 Cheeger-Type Bounds

There is a naturally motivated relationship between the expansion properties of a graph, and the eigenvalues of the Laplacian. Cheeger [Che70] proved a lower bound on the smallest positive eigenvalue for the Laplacian on Riemannian manifolds. Here we give two discrete versions of that theorem which hold for reversible Markov chains.

As before, let P be a reversible ergodic chain on V with stationary distribution π . If S is any subset of states, let $\bar{S} = V - S$ and

$$C(S) = \{e \in E \mid e = \{x, y\}, x \in S, y \in \bar{S}\}$$

denote the set edges crossing the "cut" between S and \bar{S} and

$$F(C(S)) = \sum_{e \in C(S)} F(e).$$

Define

$$h = \min_{S: \pi(S) \leq 1/2} \frac{F(C(S))}{\pi(S)}.$$

The notation h follows Cheeger, but following Sinclair and Jerrum, who proved the next theorem, we call this the conductance of P . Intuitively, this quantity is a measure of the ability of the chain to admit "flow" $F(C(S))$ of probability mass out of any set S , adjusted by the weight of S in the stationary distribution. If h is high, there are no "bottlenecks" in the flow, and one expects that convergence will then be rapid. The following theorem confirms this intuition.

In the case that P is the random walk on a d -regular graph, note that the quantity h reduces to

$$h = \min_{S: |S| \leq |V|/2} \frac{|C(S)|}{d|S|}.$$

Theorem 1.16 (Conductance Bound) [SJ87]

Let P be an ergodic time-reversible Markov chain, let L be the Laplacian for P , and let μ_1 be the smallest positive eigenvalue of L . If h is the conductance of P defined above, then μ_1 satisfies

$$\frac{h^2}{2} \leq \mu_1 \leq 2h.$$

There is another version, based on a 'vertex' expansion quantity called magnification, which we denote c . In some cases it will give better bounds if we can avoid converting to the edge-based version when a magnification bound is known.

For a given set of vertices $S \subset V$, let $\text{Nbd}(S)$ denote the set of nodes that are neighbors of nodes in S , but are not in S :

$$\text{Nbd}(S) = \{y \mid y \notin S, (\exists z \in S)[(z, y) \in E]\}.$$

Define

$$c = \min_{S: |S| \leq |V|/2} \frac{|\text{Nbd}(S)|}{|S|}.$$

Call a graph that has magnification c a c -magnifier.

Theorem 1.17 (Magnification Bound) [Alo86] Let P be an ergodic symmetric chain with underlying graph G_P , and let $p = \min_{\{x,y\} \in E} P_{xy}$ and $p' = \max_{\{x,y\} \in E} P_{xy}$. Let L be the Laplacian of P , and μ_1 be its smallest positive eigenvalue. If G_P is a c -magnifier then μ_1 satisfies

$$\frac{pc^2}{6} \leq \frac{pc^2}{4+2c^2} \leq \mu_1 \leq \frac{p'c}{2(1-c)}.$$

Proof: Apply Theorem 1.15 together with the lower bound and upper bounds for $\nu_1(Q(G_P))$ provided by Lemma 2.4 and Theorem 2.5 of Alon [Alo86]. \blacksquare

A family of graphs $\mathcal{G} = \{G_n\}$ where $|G_n| = n$, is a family of (d, c) -magnifiers if for all n , G_n is d -regular (d constant in n) and magnification at least c (again constant in n). It is an immediate consequence of the preceding theorem that a family of d -regular graphs is a magnifying family precisely when $\nu_1(G_n)$ is bounded away from 0 as n increases, (equivalently if the same holds for μ_1 for the natural random walk on G_n). By Theorem 1.8, the random walk on G_n has a convergence guarantee that is $O(\ln n + \ln(1/\epsilon))$. It is easy to see, that as a function of n , this is the best possible; no family of d -regular graphs (with d fixed) can have a convergence guarantee that is $o(\ln n)$, as n grows (taking any fixed $\epsilon < 1$). For, if the walk starts on a vertex v_0 in G_n , so that π_0 is a point-mass on v_0 , and if $k = f(n) = o(\log_d n)$ steps are taken, then the distribution π_k is supported on at most $d^k = o(n)$ vertices. Since G_n is regular, the stationary distribution is uniform, thus the variation distance $\|\pi_k - \pi\|$ will approach 1 as n grows.

For large n , almost all d -regular graphs of size n are good magnifiers [Alo86], so for most large d -regular graphs G , the random walk on G converges rapidly.

Although we generally know that we are not dealing with a magnifying family, bounds on the quantities h and c will still yield bounds on the convergence rate. Bounding h or c means solving what is called an "isoperimetric problem," the graphical version of: 'How much volume can one enclose with given perimeter?' Equivalently, one can ask how large a perimeter is necessary to enclose a given volume. The Greeks knew that in the plane, the circle had smallest perimeter for given enclosed area, but probably had no proof. Similar results hold in Euclidean space of higher dimensions [Ban80]. The problem becomes more interesting when the question is asked in a bounded domain, where the boundary of the domain is not counted in the perimeter of a region. This is the sort of problem we have, except that here the domain is a graph, volume is the number of nodes, and the perimeter of a region is the number of edges or nodes at the region's boundary.

Remark: The conductance version was first proved by Sinclair and Jerrum [SJ87]. Diaconis and Stroock [DS89] give a simpler proof. The vertex-based version of Cheeger's inequality was proved earlier by Alon [Alo86] using the Max-Flow/Min-Cut theorem. Alon's goals were in the reverse direction, to give lower bounds on expansion properties using knowledge of the eigenvalue. \square

1.5.3 Canonical Path Arguments

Canonical path arguments give a technique for lower-bounding the Cheeger quantity h . Systems of paths are also used in the Poincaré-type bounds presented later.

Let P be a reversible ergodic chain. For each z and y choose some canonical path γ_{zy} from z to y . Call this system of paths Γ . Define the (weighted) covering number of Γ

$$\eta(\Gamma) = \max_{e \in E} \frac{1}{F(e)} \sum_{\gamma_{zy} \ni e} \pi(z)\pi(y), \quad (1.9)$$

where the maximum is over all possible transitions e of the chain, and the sum is over paths γ_{zy} that contain e . This is a measure of how heavily the system of paths uses any one edge.

Theorem 1.18 (Conductance in Terms of Paths) *If P has a system Γ of canonical paths with $\eta = \eta(\Gamma)$, then the conductance h of P satisfies:*

$$h \geq \frac{1}{2\eta}.$$

Proof: Let $S \subseteq V$ with $\pi(S) \leq \frac{1}{2}$. Then

$$\begin{aligned} \pi(S)\pi(\bar{S}) &= \sum_{z \in S, y \in \bar{S}} \pi(z)\pi(y) \\ &\leq \sum_{z \in S, y \in \bar{S}} \pi(z)\pi(y) \sum_{e \in (\gamma_{zy} \cap C(S))} \frac{F(e)}{F(e)} \end{aligned} \quad (1.10)$$

$$= \sum_{e \in C(S)} F(e) \left(\frac{1}{F(e)} \sum_{z \in S, y \in \bar{S} \wedge \gamma_{zy} \ni e} \pi(z)\pi(y) \right) \quad (1.11)$$

$$\leq \eta \sum_{e \in C(S)} F(e) \quad (1.12)$$

$$= \eta F(C(S)). \quad (1.13)$$

In line 1.10 we have used the fact that for any $z \in S$ and $y \in \bar{S}$, the path γ_{zy} must cross the cut, and so contain an edge in $C(S)$. This insures we are multiplying by a factor that is at least 1. Now using the fact that $\pi(S) \leq \frac{1}{2}$, so that $\pi(\bar{S}) \geq \frac{1}{2}$,

$$\frac{F(C(S))}{\pi(S)} \geq \frac{1}{2\eta}.$$

■

The theorem has an easy combinatorial interpretation in the case that P is a random walk on a regular graph. To illustrate, let G be a d -regular connected non-bipartite graph, and let P be a random walk on G . Recall that then

$$h = \min_{S: |S| \leq |V|/2} \frac{|C(S)|}{d|S|}.$$

The problem of lower-bounding h reduces to that of giving a lower bound for the ratio $|C(S)|/|S|$. Suppose that we have a system of paths Γ such that no edge appears more than b times (total over all paths). We call b the unweighted covering number of Γ .

Now consider an arbitrary subset of vertices $S \subseteq V$. There are exactly $|S| \times |\bar{S}|$ paths γ_{vw} that go from S to \bar{S} . Since no edge e from S to \bar{S} appears more than b times over all these paths, there must be at least $|S||\bar{S}|/b$ edges that cross from S to \bar{S} . Thus $C(S) \geq |S||\bar{S}|/b$. Combining this with the supposition that $|S| \leq |V|/2$ (hence $|\bar{S}| \geq |V|/2$), we get

$$h \geq \frac{|V|}{2db}.$$

If we calculate η for this chain, we get $\eta = db/|V|$, and thus the same $h \geq \frac{|V|}{2db}$. For general graphs, we have the following.

Theorem 1.19 *If P is the natural random walk on a connected non-bipartite graph G with maximum degree d , and G admits a system of canonical paths in which no edge appears more than b times then*

$$h \geq \frac{\sum_v \deg v}{2d^2b}.$$

Proof: Since the stationary distribution is

$$\pi(v) = \frac{\deg v}{\sum_v \deg v},$$

we have for every $e \in \tilde{E}$,

$$F(e) = \frac{1}{\sum_v \deg v}.$$

Plugging in the bounds of the hypothesis gives

$$\eta \leq \frac{d^2 b}{\sum_v \deg v},$$

yielding the desired bound. \blacksquare

Plugging this into Theorem 1.16, gives the following.

Corollary 1.20 *With the hypothesis of the preceding theorem,*

$$\mu_1 \geq \frac{1}{2} \left(\frac{\sum_v \deg v}{2d^2 b} \right)^2.$$

So if we can define a system of canonical paths where no edge appears in too many paths, we will get a small b , yielding a large Cheeger value h , yielding a good lower bound on μ_1 for the Laplacian. Any system of canonical paths is sufficient to give an immediate bound; see Theorem 1.27. We give two simple examples here.

Example 1.21 (Line) Let G be the graph on the set $V = [n]$ where there is an edge joining every pair of vertices whose absolute difference is 1, and a self loop edge on each endpoint. We call this the n -segment. The self-loop at each endpoint makes the graph regular and the walk aperiodic. There is a unique shortest path between every pair of vertices x and y . Choose γ_{xy} to be this path. Consider any edge $e = \{x, x+1\}$. In a given direction, without loss of generality say left-to-right, e is traversed by those paths $\gamma_{v,w}$ for which $v \leq x$ and $w \geq x+1$. That is, the directed edge $(x, x+1)$ is contained in $x(n-x)$ paths. This is at most $n^2/4$, so no edge is contained in more than $b = n^2/4$ paths. We conclude that $h \geq 1/n$. Moreover, for even n , cutting the line in the middle so that $S = [n/2]$, we get just this value of h . So this is an optimal conductance bound. This gives $\mu_1 \geq 1/(2n^2)$, which is correct within a constant factor, the actual value being $1 - \cos(\pi/n) \approx \pi^2/n^2$. Using this, for the strongly aperiodic walk we get a bound of $\lambda_* \leq 1 - 1/n^2$. Thus $T_P(\epsilon) = n^2 \ln(2n/\epsilon)$ is a convergence guarantee for the strongly aperiodic walk. This is too large by about a $\ln n$ factor. The error is due to the fact that the bound is based only on the second-largest eigenvalue. A coupling argument (later) shows $O(n^2)$ convergence, and an $O(n^2)$ bound using the full spectrum is given in Chapter 2. \square

Example 1.22 (Hypercube) Let $V = \{0, 1\}^d$, and let G be the graph of the d -dimensional hypercube, with a self-loop edge at each vertex. This is the graph on V with an edge between every two binary d -tuples that differ in at most one coordinate. For any two vertices x and y define the canonical path from x to y as that path which brings x into agreement with y by "correcting" each

coordinate from left to right in the binary representation. Thus the canonical path from $(0, 1, 1, 0)$ to $(1, 0, 1, 1)$ in the 4-dimensional case is: $(0, 1, 1, 0)$ to $(1, 1, 1, 0)$ to $(1, 0, 1, 0)$ to $(1, 0, 1, 1)$.

No directed edge e is used by more than 2^{d-1} paths. To see this, first note that any edge e in the path from $x = (x_1, x_2, \dots, x_d)$ to $y = (y_1, y_2, \dots, y_d)$ joins two points of the form

$$(y_1, y_2, \dots, y_{i-1}, x_i, x_{i+1}, x_{i+2}, \dots, x_d) \text{ and } (y_1, y_2, \dots, y_{i-1}, y_i, x_{i+1}, x_{i+2}, \dots, x_d).$$

The i th coordinate is the one “corrected” by this step of the path. Now if we were given the edge e and the binary $(d-1)$ -tuple $(x_1, x_2, \dots, x_{i-1}, y_{i+1}, y_{i+2}, \dots, y_d)$ we could determine all of the coordinates of x and of y , and hence we could reconstruct the whole path γ_{xy} . Thus we can encode the set of paths through e by elements of $\{0, 1\}^{d-1}$. More formally, there is an injection from $\{e\} \times \{0, 1\}^{d-1}$ into $\{\gamma_{xy}\}$. It follows that no more than 2^{d-1} paths use a given directed edge e .

Consequently, $h \geq 1/(d+1)$. This gives a bound of

$$\mu_1 \geq 1/(2(d+1)^2),$$

which implies that

$$T_P(\epsilon) = 4(d+1)^3 \ln(2/\epsilon)$$

is a convergence guarantee for the strongly aperiodic walk. This is an order of $d^2/\ln d$ from the “right” answer, as we will discover later by obtaining the complete spectrum.

Although the eigenvalue bound is incorrect, the lower bound on the conductance h obtained by this argument is optimal; there are cuts $S \subset V$ attaining this value of h . For example, take S to be the subcube of points with a 0 in the first coordinate. It should be noted that an argument based on c (magnification) does not do asymptotically better. (The relevant minimal cut for determining c for d even is a Hamming ball centered at $(0, 0, \dots, 0)$ and containing half the vertices, and this also gives the same $\mu_1 = \Omega(1/d^2)$ bound. For proofs see [Bol86, Chapters 5 and 16].) \square

There is a second idea hidden in the last example. It is a method of bounding the breadth b , for a given system of canonical paths, and this aspect of the argument is due to Jerrum and Sinclair [JS88]. Suppose we can “encode” the set of paths that use any edge e by elements of a set B_e so that given e , any particular path γ_{xy} that contains e can be reconstructed given the additional information of an element from B_e . Then at most $b = \max_e |B_e|$ use a given edge. In the previous example, paths containing a given edge were encoded using the $(d-1)$ -tuple, and $B_e = \{0, 1\}^{d-1}$, $|B_e| = 2^{d-1}$ for all e .

1.5.4 Poincaré-type Bounds

Poincaré-type bounds are another type of bound based on canonical paths. These also rely on the careful choice of a system of canonical paths having small covering number.

For a system of canonical paths Γ for P , define the quantity

$$K(\Gamma) = \max_e \frac{1}{F(e)} \sum_{\gamma_{xy} \ni e} |\gamma_{xy}| \pi(x) \pi(y). \quad (1.14)$$

Here again, the maximum is over transitions e , the sum is over paths containing e , and $|\gamma_{xy}|$ denotes the length of the path.

Theorem 1.23 (Poincaré Bound) [DS89] *Let P be a reversible ergodic chain, and Γ a system of canonical paths for P , then*

$$\mu_1 \geq \frac{1}{K(\Gamma)}.$$

Proof: We will lower-bound the Rayleigh quotient in 1.7. Using the fact that $(\phi, 1)_\pi = 0$, write

$$(\phi, \phi)_\pi = \frac{1}{2} \sum_{x,y} (\phi(x) - \phi(y))^2 \pi(x) \pi(y).$$

Now write the difference $(\phi(x) - \phi(y))$ as a telescoping sum along the path γ_{xy} :

$$(\phi(x) - \phi(y)) = \sum_{e \in \gamma_{x,y}} \nabla \phi(e).$$

Combining these we proceed

$$\begin{aligned} (\phi, \phi)_\pi &= \frac{1}{2} \sum_{x,y} \left(\sum_{e \in \gamma_{x,y}} \nabla \phi(e) \right)^2 \pi(x) \pi(y) \\ &\leq \frac{1}{2} \sum_{x,y} |\gamma_{xy}| \sum_{e \in \gamma_{x,y}} (\nabla \phi(e))^2 \pi(x) \pi(y) \\ &= \frac{1}{2} \sum_{e \in E} (\nabla \phi(e))^2 \sum_{\gamma_{x,y} \ni e} |\gamma_{xy}| \pi(x) \pi(y) \\ &\leq K(\Gamma) \frac{1}{2} \sum_{e \in E} (\nabla \phi(e))^2 F(e) \\ &= K(\Gamma) (\phi L, \phi)_\pi. \end{aligned} \quad (1.15)$$

and thus the quotient in 1.7 is bounded below by $1/K(\Gamma)$, giving the stated result. The inequality in 1.15 is Cauchy-Schwarz. ■

The following is a simple corollary, obtained by noting that if the longest path $\gamma_{xy} \in \Gamma$ has length ℓ , then $K(\Gamma)$ is related to the weighted covering number by $K(\Gamma) \leq \ell \eta(\Gamma)$.

Corollary 1.24 *Let P be an reversible ergodic chain admitting a system Γ of canonical paths each of length at most ℓ , and let $\eta = \eta(\Gamma)$. Then*

$$\mu_1 \geq \frac{1}{\ell \eta}.$$

This yields the following result in the graphs realm.

Corollary 1.25 *Let $G = (V, E)$ be a connected non-bipartite graph with maximum degree d , admitting a system of canonical paths such that no edge appears more than b times over all paths, and such that each path has length at most ℓ . Then the random walk on G has*

$$\mu_1 \geq \frac{\sum_v \deg v}{\ell d^2 b}.$$

Proof: Using the same reasoning as in Theorem 1.19 to bound η , apply the previous theorem. \square

Remark: The ideas behind the discrete Poincaré bounds are due to Diaconis and Stroock [DS89]. The proof above is translated into our language. Unlike in Jerrum and Sinclair's Cheeger-type inequality, the paths here enter the proof in a fundamental way; a key step is to write $\phi(x) - \phi(y)$ as a telescoping sum of the gradient along the path from x to y . These bounds often give better results than the Cheeger-type bounds, based on the same system of canonical paths. Note that there is no squaring involved here, as arises in the Cheeger bound. Using a given system of canonical paths, and the versions of the theorems involving η , the Poincaré-type bound will give a better result than the Cheeger-type bound if and only if $\ell < 8\eta$, and in practice, this seems almost always to be the case. However, it may happen that one has a bound on c or h without knowing a good system of canonical paths, so that a Cheeger-type bound may be available, though no Poincaré-type bound is evident. This is the case with some of our later bounds for the eigenvalues of Cayley graphs. \square

Example 1.26 (Hypercube) Take, again, the hypercube with self-loops. Using the same system of canonical paths, as in Example 1.22, we have $b = 2^{d-1}$, $\ell = d$, and the graph is $(d+1)$ -regular. So the Poincaré bound for the walk on the hypercube gives $\mu_1 \geq \frac{2}{d(d+1)}$. This is better than the conductance bound by a constant factor, but is still of the same order. It is still a factor of d from the right value $\mu_1 = \frac{2}{d+1}$. We should note that the Poincaré bound can be used in another way to get a bound of the form $\mu_1 = \Omega(1/d)$. Project the walk on the hypercube to the line of length d by recording the distance from the origin. Given the distance from the starting point, the walk on the hypercube is uniformly distributed on the points at that distance. Carefully calculating with the resulting non-uniform edge probabilities it is possible to get a bound on μ_1 that is accurate to within a constant factor. See [DS89]. \square

The ideas of the previous sections can be combined to give a general bound, which is of limited interest to us here, but can be applied to such problems as covering times and universal traversal sequences. (See [BK88] [AKL⁺79]).

Theorem 1.27 (General Bounds) *Let G be a connected non-bipartite graph with maximum degree d and diameter Δ . Let P be the natural random walk on G . Then the smallest positive eigenvalue*

μ_1 of $L = I - P$ satisfies

$$\mu_1 \geq \max \left\{ \frac{\sum_v \deg v}{d^2 \Delta |V|^2}, \left(\frac{\sum_v \deg v}{d^2 |V|^2} \right)^2 \right\}.$$

Proof: Use any system of shortest paths as the canonical paths. Then $\ell = \Delta$, and $b \leq |V|^2$. Apply the Cheeger and Poincaré bounds. \blacksquare

1.6 Probabilistic Bounds

We have been talking exclusively about spectral bounds on the convergence rate. There are also some purely probabilistic techniques for bounding the convergence rate. These techniques are attractive because they involve only elementary and ‘algorithmic’ reasoning, and give a more intuitive description of time to stationarity. However, they also have drawbacks. First, they seem to require a certain cleverness or the use of rather special structure in order to yield good bounds. Second, in our applications it is often helpful to know the spectrum, and it is not known how to recover this information accurately from the probabilistic bounds.

Let P be a Markov chain on a finite state space V . Let S^∞ denote the set of all infinite sequences from the state space S . For an element $\sigma \in S^\infty$, let $\sigma[1:k]$ denote the initial segment consisting of the first k elements of the sequence. A stopping time T for P is a function $T: S^\infty \rightarrow N \cup \{\infty\}$, such that if $T(\sigma) = k$ and $\sigma'[1:k] = \sigma[1:k]$ then $T(\sigma') = k$ as well. In other words, if T assigns a value k to a given sequence of states σ then it does so to all sequences sharing the same initial segment up to k . Intuitively, T determines a value k that only depends on observing the sequence up to time k . We regard T as a statistic on S^∞ under the infinite product measure given by the π_k ’s, and we may then speak of the distribution of the random variable T .

The techniques discussed in the following two sections are based on the construction of stopping time random variables T with the property that the variation distance $\|\pi_k - \pi\|$ of the k th-step distribution of the walk from the stationary distribution is bounded by $\Pr\{T > k\}$, the mass in the “tail” of the distribution of T after k .

1.6.1 Couplings

Let Z be an ergodic Markov chain with state space V , and transition matrix P . Let π_0 be its initial state distribution, let π_k denote the distribution of its state at time k (i.e., $\pi_k = \pi_0 P^k$), and let π be its unique stationary distribution.

A (Markovian) coupling (X, Y) for Z is a Markov process (X, Y) on $V \times V$ together with a stopping time T for (X, Y) such that:

(C1) The projections X and Y have transition matrix P , identical to Z ’s, on V .

(C2) Y is started in a state chosen according to the stationary distribution, $\pi_0^Y = \pi$. X is started with the initial distribution of Z , $\pi_0^X = \pi_0$.

(C3) T is a stopping time for (X, Y) that also satisfies: if $T = t$ then $X_k = Y_k$ for all $k \geq t$. The coupling is proper if T is finite with probability 1. We will generally use 'coupling' to mean 'proper coupling.' If it occurs that $T = t$, the coupling is said to have succeeded at time t .

Note that in the process (X, Y) , the projections X and Y need not be independent, and in general won't be. The requirement that the pair-process (X, Y) be Markovian can be removed, as long as the projections remain Markovian like Z . Thorisson [Tho86] has noted that the requirement that X and Y meet and remain identical from time T onwards can be weakened to the requirement that X and Y become and remain identically distributed from that time on, without substantially changing the proof of the following theorem. Also, one can always adjust the coupling to ensure that they remain together once they meet.

Theorem 1.28 (Coupling Inequality) *Let Z be a Markov process on V with stationary distribution π , and let (X, Y) be a coupling for Z with coupling time T . Let π_k be the distribution of the state of Z after k steps. Then for all $k \geq 0$, we have*

$$\|\pi_k - \pi\| \leq \Pr\{T > k\}.$$

Proof: Let A be any subset of S . Let z_k , x_k and y_k denote the states of Z , X , and Y , respectively, after the k th step. Then we have

$$\begin{aligned} |\pi_k(A) - \pi(A)| &= |\Pr\{z_k \in A\} - \pi(A)| \\ &= |\Pr\{z_k \in A\} - \Pr\{y_k \in A\}| \\ &= |\Pr\{z_k \in A \wedge k \geq T\} + \Pr\{z_k \in A \wedge T > k\} \\ &\quad - (\Pr\{y_k \in A \wedge k \geq T\} + \Pr\{y_k \in A \wedge T > k\})| \\ &= |\Pr\{z_k \in A \wedge T > k\} - \Pr\{y_k \in A \wedge T > k\}| \\ &\leq |\max(\Pr\{z_k \in A \wedge T > k\}, \Pr\{y_k \in A \wedge T > k\})| \\ &\leq \Pr\{T > k\}. \end{aligned}$$

In the second line we have used the fact that the distribution of X and Z are identical, and the fact that Y is started, and thereby remains, in the stationary distribution. In the fourth line we use the fact that X and Y coincide from time T onwards. \blacksquare

Couplings have their origins in proofs of ergodic theorems for more general stochastic processes, rather than in convergence rate proofs for Markov chains. Showing the existence of a stationary distribution for a process, and displaying a coupling with $\Pr\{T > k\}$ going to zero with increasing k , gives an ergodic theorem. (See Griffeath [Gri78] for a survey of coupling methods for Markov processes on general state spaces.) For our purposes, we want our processes to have an easily analysable coupling time that closely matches the actual rate of convergence.

Griffeath [Gri75] shows that there always exists a maximal coupling, one in which

$$\|\pi_k - \pi\| = \Pr\{T > k\}$$

for all k , thus achieving equality in the Coupling Inequality. Unfortunately, constructing such a coupling usually requires more knowledge of the chain than we have. However, we can sometimes make use of the fact that such couplings exist, without actually constructing one. (See, e.g. Theorem 2.8.)

Here are some simple examples of couplings. We give some new coupling bounds in Chapter 2.

Example 1.29 (Line) We consider again the random walk P on the n -segment. (See Example 1.21.) Our strategy for (X, Y) is to run the chain X normally like P , and to mimic X with Y . Let X start out at the left end-point and Y at a uniformly chosen point. Choose a move for X normally. If X moves right, move Y right (or at the right endpoint take the self-loop). If X moves left, move Y left (or at the left endpoint take the self-loop). This clearly makes X and Y each behave like P . Notice that when the X and Y meet, they remain together. Let T be the time they first meet. This clearly gives a coupling. To bound the tails on T , note that by following this strategy for (X, Y) , X can never move to the right of Y . So that in particular, by the first time Y reaches the left boundary, X and Y must have met. It is a well-known fact that the random walk on such a segment hits the left boundary in $O(n^2)$ steps. See [Fel70, Vol I., Ch. XIV] for a detailed analysis of this time. Note that this is the 'right answer,' as we show in Chapter 2. \square

Example 1.30 (Convergence of a Fixed Finite Chain) We can get a proof of the convergence portion of the basic ergodic theorem for Markov chains as follows. Let P be an ergodic Markov chain. We will show that if π is any stationary distribution of P , then the chain started with any distribution converges to π . This will also imply uniqueness of π . Recall, P is ergodic means that there is a k_0 such that $P_{xy}^{(k)} > 0$ for all $k \geq k_0$. Let $p > 0$ be the smallest entry in P^{k_0} . Now consider the following coupling (X, Y) . Choose an initial value for X according to any distribution, and choose an initial value for Y according to π . Since π is stationary Y will remain so distributed. Now run both (X, Y) , each independently according to P . Let T be the first time they meet. Now after k_0 steps, the point X is in some state x , and the probability that Y is in the same state x is at least p . It follows from the Coupling Inequality that $\|\pi_{k_0} - \pi\| \leq \Pr\{T > k_0\} \leq (1-p)^k$. Thus π_k converges to π and immediately "exponentially fast." This example is somewhat misleading on two accounts. Here the chain is fixed; p will not in general be constant in a family of chains for which the vertex set grows. Furthermore, spectral analysis is usually required anyway to get any reasonable handle on p (or anything like it). \square

1.6.2 Strong Stationary Times

A strong stationary time for a Markov chain $\{X_k\}$ with stationary distribution π is a stopping time T such that

$$\Pr\{X_k \in A \mid T \leq k\} = \pi(A).$$

When π is the uniform distribution on the state space we also call T a **strong uniform time**. For background and additional material see [AD86] and [DF88].

Essentially the same proof as that of the Coupling Inequality will verify that a strong stationary time satisfies the same inequality

$$\|\pi_k - \pi\| \leq \Pr\{T > k\}.$$

Indeed strong stationary times can be considered a special type of coupling.

In the context of strong uniform times, it is somewhat more natural to measure distance between distributions by the **separation**, defined by

$$\text{sep}(P, Q) = \max_{s \in S} \frac{Q(s) - P(s)}{Q(s)}.$$

(Separation is discussed further in Appendix A.) This is because strong stationary times satisfy the stronger inequality

$$\|\pi_k - \pi\| \leq \text{sep}(\pi_k, \pi) \leq \Pr\{T > k\}.$$

Just as there always exists a maximal coupling, there is always a maximal strong stationary time, one which achieves equality here. However, since we may have $\|\pi_k - \pi\| < \text{sep}(\pi_k, \pi)$, strong uniform times, in general, cannot be maximal couplings. (However, when the Markov chain is a random walk on a group, one can show that if the variation gets small after k steps, the separation is small after at most roughly $2k$. So, roughly speaking, a maximal strong stationary (uniform) time for a random walk on a group can only be off by about a factor of two. For a precise statement and proof see [AD86, Proposition (5.13), p. 15].)

Example 1.31 (Top-in-at-Random Shuffle) Suppose a deck of n cards is shuffled by the following procedure. Take the top card, and place it in a uniform random position in the deck (possibly back on top). This defines a doubly stochastic Markov chain. How long does it take to converge to uniformity? Let T be the first time the card, call it B , that was originally on the bottom comes to the top and has just been re-inserted. T is a strong uniform time. This can be proved by induction as follows. Note that the card B rises only by having a card inserted below it. The relative order of the cards below B is uniform and each time a new card is inserted, this property is preserved. Thus, when finally B reaches the top and is re-inserted, all of the cards are in random relative order. To analyze T , note that $T = \sum_{1 \leq i \leq n} T_i$, where T_i is a geometric random variable with parameter i/n . This tells us that $E[T] = nH_n = n(\ln n + \gamma + o(1))$. The tail after this point goes down exponentially and this can be seen easily using a coupon collecting argument. (See Chapter 2.) \square

Chapter 2

Urn Models and Group Actions

Urn transfer models are processes in which the states are configurations of balls in urns, and whose evolution proceeds according to Markovian transfer rules for the balls amongst the urns. Such processes often arise as models of diffusion and as particle models in statistical mechanics. Various applications are described in [JK77].

While it is easily seen that any finite-state Markov chain can be viewed in a trivial way as a single-ball urn model, one gets a more natural class of models by insisting that the transfer rules have certain symmetries. In particular, we will require that they correspond to the action of the generators of some group.

In this chapter, we give some upper bounds on the time to stationarity of such processes. We first present methods of dealing with two well-known urn models. We then give general diameter-based bounds on the Laplacian eigenvalues and rate of convergence for certain Markov chains based on group actions.

2.1 Ehrenfest-type Models

2.1.1 Arbitrary Steps

Consider a system of n labelled balls distributed amongst m labelled urns, where the configuration of the system evolves in discrete time according to the following rule. At each time step, one of the n balls is chosen uniformly at random, and this ball ‘jumps’ from the urn it is in to a new location chosen uniformly at random amongst the urns (including the one in which the ball started). We call this chain the Ehrenfest- (n, m) process after the physicists who proposed the $m = 2$ case to model certain bi-valued properties of particles [EE07, Fei70].

This process is a symmetric ergodic Markov chain. By associating to each ball the number of the urn it is in, we obtain a correspondence of the states of the system to ordered n -tuples where

each coordinate is drawn from the set $[m]$. Using the standard Cartesian product notation for sets, we denote this $V = [m]^n$. The stationary distribution is uniform on the state space.

Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be graphs (in our general sense) with $|V_1| = n_1$ and $|V_2| = n_2$. Let A_1 and A_2 be their adjacency matrices. The Cartesian product, $G = G_1 \times G_2$, is the graph whose vertex set is $V_1 \times V_2$ and whose adjacency matrix is given by

$$A_1 \otimes I_{n_2} + I_{n_1} \otimes A_2,$$

where I_n denotes the identity matrix of order n and \otimes denotes the Kronecker product. (So $A_1 \otimes I_{n_2}$ is the $n_1 \times n_2$ matrix obtained by placing n_2 'copies' of A_1 along the diagonal, and $I_{n_1} \otimes A_2$ is obtained from A_2 by replacing a_{ij} by a diagonal block of n_1 copies of a_{ij} .) Informally, the product graph may be constructed by first making a copy G_v of G_1 for each vertex in $v \in V_2$, and joining each corresponding pair of vertices in G_v and G_w , if (v, w) is an edge in G_2 . Multiple edges and self-loops are copied with their corresponding multiplicities. The Cartesian product is an associative operation on graphs.

The Ehrenfest- (n, m) process may be viewed as the natural random walk on the nm -regular graph obtained by adding n self-loops to each vertex of K_m^n , the Cartesian product of n cliques of size m . When $m = 2$, this is the n -dimensional hypercube with n self-loops at each vertex. Diaconis [Dia88] gives arguments to deal with this case. In this section we generalize to larger m using a different method.

Lemma 2.1 (Laplacian Spectrum of Cartesian Products) *Let $H = (V_H, E_H)$ and $K = (V_K, E_K)$ be two graphs. Let $G = H \times K$ be their Cartesian product. Then the eigenvalues of $Q(G)$ are the elements of the multiset of all possible sums $\nu_H + \nu_K$, where ν_H is an eigenvalue of $Q(H)$ and ν_K is an eigenvalue of $Q(K)$. In particular $\nu_1(G) = \min\{\nu_1(H), \nu_1(K)\}$.*

Proof: The statement follows easily from the structure of the matrix $Q = D - A$ for the product graph G in terms of the corresponding matrices for the components G_1 and G_2 . Again, let $n_1 = |V_1|$, $n_2 = |V_2|$, and let I_n denote the identity matrix of order n . One can verify that

$$Q(G) = Q(G_1) \otimes I_{n_2} + I_{n_1} \otimes Q(G_2).$$

The result then follows from the well-known theorem of Kronecker that for square matrices A and B , of order a and b respectively, the eigenvalues of the product $A \otimes B$ are the ab possible products of eigenvalues, one from A and one from B . (See e.g. [MM64, p. 24, Prop. 2.15.11].) \blacksquare

Theorem 2.2 (Eigenvalues of the Ehrenfest Process) *Let P be the Ehrenfest- (n, m) process. Let $L = I - P$ be the associated Laplacian. The eigenvalues of L are the values*

$$\frac{j}{n}, \quad 0 \leq j \leq n, \quad \text{with multiplicity } \binom{n}{j} (m-1)^j.$$

Proof: First consider K_m the m -clique. For this graph we have $Q(K_m) = mI - J$, where J is the $m \times m$ matrix with unit entries throughout. The characteristic polynomial is $\det(Q - zI) = z(n - z)^{n-1}$. Thus its eigenvalues are 0 and m , where m occurs with multiplicity $m - 1$. It follows from the preceding lemma that the spectrum of $Q(K_m^n)$ is the set of values m_j , $0 \leq j \leq n$, where the value m_j occurs with multiplicity $\binom{n}{j}(m-1)^j$, and that the eigenvalues of $L(P) = \frac{1}{nm}Q(G)$ are $\frac{m_j}{nm} = \frac{j}{n}$ with the same multiplicities. \blacksquare

Corollary 2.3 (Full-spectrum bound for the Ehrenfest process) *For any initial distribution, let π_k be the k th-step distribution of the Ehrenfest- (n, m) process. If $k > \frac{n}{2}(\ln n + \ln m + c)$ for any $c > 0$, we have*

$$\|\pi_k - U\|^2 \leq \frac{1}{4}(e^{e^{-c}} - 1).$$

Proof: Apply the previous result and Theorem 1.9. Taking $k > \frac{n}{2}(\ln n + \ln m + c)$, one gets

$$\begin{aligned} \|\pi_k - U\|^2 &\leq \frac{1}{4} \sum_{j=1}^n \binom{n}{j} (m-1)^j \left(1 - \frac{j}{n}\right)^{2k} \\ &\leq \frac{1}{4} \sum_{j=1}^n \frac{(nm)^j}{j!} e^{-2kj/n} \\ &= \frac{1}{4} \sum_{j=1}^n \frac{e^{-jc}}{j!} \\ &\leq \frac{1}{4}(e^{e^{-c}} - 1). \end{aligned}$$

Here, in moving to the second line we have used the fact that $(1 - x) \leq e^{-x}$, and in the last line, we use the fact that the sum is an initial segment of the series expansion for $e^{e^{-c}} - 1$. \blacksquare

For $m = 2$ the bound is essentially tight [Dia88]. For n fixed and m growing the bounds obtained by the preceding technique grow with m . In actuality, however, the time to stationarity is independent of m , as is shown in the next theorem, by an easy strong uniform time bound. Knowledge of the spectrum, however, is useful in applications, such as in Chapter 3.

Theorem 2.4 (Uniform Time for Ehrenfest- (n, m)) *Consider the Ehrenfest- (n, m) process, and let T be the first time each ball has been chosen at least once to move. Then T is a strong uniform time.*

Proof: Mark the coordinate corresponding to a ball immediately after it is first chosen to move. The value in the first marked coordinate, upon being marked, uniformly distributed over its m possibilities by the definition of the process. Similarly, as a simple induction shows, the restriction of the state to the marked coordinates is always uniformly distributed, and independent of the time

the set was so marked. It follows that the first time all coordinates are marked is a strong uniform time for the process. \square

The analysis of the above strong uniform time is the same as that of the "coupon collector's" problem. An easy bound can be obtained as follows.

Corollary 2.5 *Let π_k be the k th-step distribution of the Ehrenfest- (m, n) process. For any initial distribution π_0 , we have*

$$\|\pi_k - U\| \leq e^{-c} \text{ whenever } k \geq n \ln n + cn.$$

Proof: Let E_{ik} denote the event that after k steps, the coordinate i remains unmarked. The event $\{T > k\}$ is clearly $\bigcup_{1 \leq i \leq n} E_{ik}$. So

$$\Pr\{T > k\} \leq \sum_{1 \leq i \leq n} \Pr\{E_{ik}\} = n \Pr\{E_{ik}\},$$

since the probability of a union of events is bounded by the sum of their probabilities, and each of the E_{ik} is equiprobable. Now, the probability that a given coordinate, say coordinate 1, is unmarked after step k is the probability it has not been chosen by step k , which is $(1 - \frac{1}{n})^k$. Whence

$$\Pr\{T > k\} \leq n \left(1 - \frac{1}{n}\right)^k = e^{ln n + k \ln(1 - 1/n)} \leq e^{ln n - k/n},$$

from which the inequality follows immediately. \square

Remark: As n grows, the coupling time T is asymptotically Poisson with mean n . For n growing we have

$$\Pr\{T \leq n \ln n + cn\} = e^{-e^{-c}} + o(1).$$

\square

Remark: The number of self loops affects the holding probabilities, and therefore the convergence rate bounds, but not significantly. The spectral argument may be used with minimal change to bound the walk on the graph for any number of self-loops, provided the graph remains regular. For $m = 2$ some self loops are necessary to avoid periodicity. For $m \geq 3$, the walk on the graph is aperiodic without the need to add self-loops. The strong uniform time argument is more sensitive to the number of self-loop edges. There is a simple coupling with the same convergence rate for the case of a single self-loop at each vertex.

Broder [Bro], and Aldous and Diaconis [AD86] have suggested similar strong uniform times for the special case of $m = 2$, when the holding probability at each state is $1/2$. The extension here to larger m is straightforward, if that is known. This result was obtained independently around the same time (1986). \square

2.1.2 Adjacent Moves

An interesting variant of the Ehrenfest model is the adjacent-move model. Here, as before, we have n labelled balls and m labelled urns. A ball is chosen uniformly; say it is currently in urn i . A value in $d \in \{-1, +1\}$ is chosen uniformly, and the ball moves to urn $i + d$ if that urn exists (i.e., $1 \leq i + d \leq m$). Otherwise it remains in the same position.

This leads to the natural random walk on a $2n$ -regular graph given by the Cartesian product of n m -segments, where each m -segment has a single self-loop at its two ends. (See our earlier Examples 1.21 and 1.29.) The stationary distribution is uniform on $V = [m]^n$.

We first consider the walk on the m -segment. The eigenvalues of this transition matrix are known classically. They are $\cos(\pi i/m)$, for $0 \leq i \leq m-1$. (See [Fel70, Vol. 1, Ch. XVI.3].) The unit eigenvalue comes from $i = 0$.

Theorem 2.6 (Full Spectrum Bound for the Line) *Let π_k be the k th-step distribution of the natural random walk on the m -segment. For any initial distribution and any real $c \geq 1$, if $k > \frac{cm^2}{\pi^2}$ then*

$$\|\pi_k - U\| \leq e^{-c}.$$

Proof: Theorem 1.9 gives us

$$\|\pi_k - U\|^2 \leq \frac{1}{4} \sum_{j=1}^m (\cos(\pi j/m))^{2k}.$$

It is not too difficult to bound this sum using properties of the cosine.

The following properties are handy. We have

$$\cos(x) \leq e^{-x^2/2} \quad \text{for } 0 < x < \pi/2. \quad (2.1)$$

To see this, let $h(x) = \ln(\cos(x)e^{x^2/2})$ and note $h(0) = \ln(1) = 0$, $h'(x) = x - \tan x \leq 0$ for $0 < x < \pi/2$. So $h(x) \leq h(0)$ in this interval, and (2.1) follows. We also use the fact that for $1 \leq i \leq m$ we have $\cos(\pi j/m) = -\cos(\pi \ell/m)$ where $\ell = m - j$, so their even powers are equal. Finally note that $\cos(\pi/2) = 0$.

Using these properties, we have

$$\begin{aligned} \frac{1}{4} \sum_{j=1}^m (\cos(\pi j/m))^{2k} &= \frac{1}{2} \sum_{j=1}^{\lfloor (m-1)/2 \rfloor} \cos(\pi j/m))^{2k} \\ &\leq \frac{1}{2} \sum_{j=1}^{\lfloor (m-1)/2 \rfloor} e^{-k\pi^2 j^2/m^2} \\ &\leq \frac{e^{-k\pi^2/m^2}}{2} \sum_{j=1}^{\infty} e^{-k\pi^2(j^2-1)/m^2} \end{aligned}$$

$$\begin{aligned} &\leq \frac{e^{-k\pi^2/m^2}}{2} \sum_{j=0}^{\infty} e^{-3k\pi^2 j/m^2} \\ &= \frac{e^{-k\pi^2/m^2}}{2(1 - e^{-3k\pi^2/m^2})}. \end{aligned}$$

The denominator is certainly greater than 1 when $k > m^2/\pi^2$. The result follows. \square

This bound is tight to within a constant factor.

Theorem 2.7 (Lower Bound for the Line) Let π_k be the k th-step distribution of the natural random walk on the m -segment, when π_0 is any point mass distribution on one vertex. For any $c > 0$, if $k \leq \frac{m^2}{16(c+1)}$ then

$$\|\pi_k - U\| \geq 1 - 2e^{-c}.$$

Proof: If $k \leq m^2/16(c+1)$ moves are taken, a standard Chernoff bound (e.g., Lemma 3.10) insures that the probability that we are at distance at least $m/4$ away from the initial point is at most $2e^{-m^2/16k} \leq 2e^{-(c+1)} \leq e^{-c}$. Rephrasing this, the total mass of π_k on that proportion of points at distance more than $m/4$ is at most e^{-c} . Let $p \leq 1/2$ be the fraction of points at distance less than $m/4$ from the initial vertex. The variation distance is at least $((1-p) - e^{-c}) + ((1-e^{-c}) - p) = 2(1-p) - 2e^{-c} \geq 1 - 2e^{-c}$. \square

Remark: Diaconis [Dia88] gives similar tight bounds for random walk on the circle Z_m , m odd. He obtains the eigenvalues for that transition matrix by harmonic analysis. They are $\cos(2\pi i/m)$ for $1 \leq i \leq m-1$. Though they are different, they give rise to the same sum as in our upper bound, and from that point we use the same argument. Our lower bound argument is different from his, but a similar technique is also suggested by Diaconis, p. 27. \square

Lemma 2.1 (about Cartesian products) tells us the whole spectrum for the Ehrenfest process as sums of cosines, but the resulting sum in the full spectrum bound is unwieldy. Lemma 2.1 also gives a bound on $\mu_1(P)$ alone which is easy to apply. But this gives an $O(n^2 m^2 \ln m)$ bound on the number of steps sufficient to get close to stationarity. We can get a better result from a coupling argument. The intuition is that the projection to any one coordinate is a walk on an m -segment. Moreover, the uniform distribution in each component induces the uniform distribution on the whole product. It will suffice to take enough steps to insure that each component gets some $\Omega(m^2)$ steps. The following theorem makes this rigorous.

Theorem 2.8 (Adjacent Move Ehrenfest Process) Let P be the adjacent-move Ehrenfest- (n, m) process, with $m \geq 4$. For any real $c \geq 1$, let $d = \frac{cm^2}{4} \ln 2n$. Then for any initial distribution and any $k > (6d - 1)n$ we have

$$\|\pi_k - U\| \leq e^{-c}.$$

For fixed c , this gives an $O(m^2 n \ln n)$ convergence guarantee.

Proof: It suffices to prove the statement for any given initial position $x \in V = [m]^n$.

We know, (see Section 1.6.1), that there exists a maximal coupling for any Markov process, i.e. a coupling that achieves equality in the coupling inequality. Write $x = (x_1, x_2, \dots, x_n)$ and for $1 \leq i \leq n$, let $C_i = (A_i, B_i)$ be a maximal coupling for the natural random walk on the m -segment, where A_i starts in x_i , and B_i starts at a point y_i chosen according to the uniform distribution on the m -segment. Let T_i be the associated stopping time random variable for the coupling C_i .

Now construct a 'full coupling' $C = (X, Y)$ for the chain P as follows. Start X at the initial vertex $x = (x_1, x_2, \dots, x_n) \in V$, and start Y at $y = (y_1, y_2, \dots, y_n)$ which, by the way the y_i were chosen, is uniformly distributed on V . At each time step choose a coordinate i in $[n]$. Make a move in the coupling process $C_i = (A_i, B_i)$, and duplicate the move of A_i in coordinate i of X , and the move of B_i in coordinate i of Y . Since C_i is a coupling for the natural random walk on the line, both A_i and B_i move according to the transition probabilities for natural random walk on the line. It is easy to see from this that X and Y move according to the transition probabilities given by P . Notice that from the point that C_i has successfully coupled coordinate i , the processes X and Y will also agree thereafter in coordinate i , since duplicating the moves of C_i must keep them coupled. Thus $C = (X, Y)$ is a coupling for P with coupling time $T = \max_i T_i$, and it remains only to bound the tail of the distribution of T .

To do this, first let $d = \pi^{-2}cm^2 \ln(2n)$, and suppose that at some time k we have taken at least d moves in coordinate i , for some particular i . The probability that X and Y do not agree by this time in their i th coordinate is precisely the probability that C_i has not coupled, in other words $\Pr\{T_i > d\}$. But because C_i is a maximal coupling, and by our earlier bounds for the walk on the m -segment, $\Pr\{T_i > d\}$ is at most $\exp(-c \ln(2n)) = e^{-c}/2n$. Thus, if at some time k we have made at least d moves in every coordinate, the probability that at least one of the C_i has not yet coupled is at most $e^{-c}/2$.

Now we will show that if we make $k \geq (6d - 1)n$ moves in the coupling C , that we will have moved at least d times in every coordinate with high probability, and thus we will have coupled with high probability.

Suppose we make k moves in the full coupling C . Consider any one particular coordinate i . Of these k moves, the number k_i of moves that take place in coordinate i (and hence coupling C_i) has the $\text{Binomial}(k, p)$ distribution, where $p = 1/n$. So, using a standard Chernoff bound (Lemma 3.10), we have

$$\Pr\{k_i < d\} = \Pr\{k_i \leq d - 1\} \leq \exp(-(kp - (d - 1))^2/4kpq),$$

where $q = 1 - p$. Noting that $pq < p = 1/n$, and then substituting the values of k , p , and d we get

$$\begin{aligned} \Pr\{k_i < d\} &\leq \exp(-(kp - (d - 1))^2/4kpq) \\ &\leq \exp(-(kp - (d - 1))^2 n/4k) \\ &= \exp(-(5d)^2 n/4k) \end{aligned}$$

$$\begin{aligned}
&= \exp(-25d^2n/4k) \\
&= \exp(-25d^2/(24d-4)) \\
&\leq \exp(-d) \\
&= \exp(-\pi^{-2}cm^2 \ln(2n)) \\
&\leq \exp(-c \ln(2n)), \text{ since } m \geq 4 > \pi \\
&\leq e^{-c}/2n.
\end{aligned}$$

So the union of events $\bigcup_{i=1}^n \{k_i < d\}$, which is to say the event that we fail to make at least d moves in some coordinate, has probability at most $e^{-c}/2$.

Thus after $k = (6d-1)n$ moves, the probability that we have not coupled in the full coupling C is at most e^{-c} , because we can only have failed to couple in one of two ways: (1) we made at least d moves in every coordinate but still failed to couple in at least one coordinate, which we've shown happens with probability at most $e^{-c}/2$, or (2) we failed to make at least d moves in some coordinate and we also failed to couple. This happens with probability at most $e^{-c}/2$. The theorem then follows from the Coupling Inequality. \square

Remark: The idea behind the coupling above can be used more generally to construct coupling bounds in Cartesian product chains from bounds on each component. The latter part of the argument is essentially an upper bound on the multiple coupon collecting problem: Suppose that you are given a sequence of coupons drawn from an infinite supply of n different types of coupons, and where the type of each successive coupon is independently and uniformly distributed over the n types. How many coupons must you be thrown in order to get at least d coupons of each type (with high probability)? In Corollary 2.5, we gave an answer to this question when $d = 1$. Several authors [NS60, ER61, Fla82] have given asymptotic answers to this and related questions for the case where d is any fixed positive integer, and n approaches infinity. To illustrate, let the random variable $K_{d,n}$ denote the number of coupons required to get d complete sets of the n different types of coupons. Erdős and Rényi [ER61] show that for every fixed positive integer d , and every real z one has

$$\lim_{n \rightarrow \infty} \Pr\{K_{d,n} < n \ln n + (d-1)n \ln \ln n + zn\} = \exp\left(-\frac{e^{-z}}{(d-1)!}\right).$$

Their proof begins to need patching near the outset if we wish to handle even $d = \Theta(\log n)$. For our problem, we in fact have $d = c\pi^{-2}m^2 \log n$, and we want to give guarantees that hold for finite n , and not only in the limit. For the specific case we consider, our argument gives

$$\Pr\{K_{d,n} > (6d-1)n\} \leq e^{-c}/2.$$

For any fixed $c > 0$, this bound is clearly optimal to within the factor 6, since for any d and n it is necessary to draw at least dn coupons to have any nonzero probability of getting d complete sets.

However, our bound may not be optimal as a function of c . Further non-asymptotic analysis of the general multiple coupon collector's problem would be interesting and useful. \square

2.2 Bernoulli-Laplace-type Models

Let n be a positive integer, $\beta = (\beta_1, \beta_2, \dots, \beta_m)$ a partition of n into m nonzero parts. We assume that $n > m \geq 2$, and $\beta_1 \geq \beta_2 \geq \dots \geq \beta_m > 0$. Consider the following discrete-time Markov process. Initially, n labelled balls are distributed into m labelled urns with balls $1 \dots \beta_1$ in urn 1, the next β_2 balls in urn 2, and so forth. Then at each time step, two urns are chosen uniformly at random. From each of the two chosen urns a ball is chosen uniformly at random, and the chosen balls are swapped; each moves to the other's urn.

We call this the Bernoulli-Laplace model with parameters n , m , and β . Daniel Bernoulli proposed it as a simple model of diffusion, and Laplace did some analysis. Similar models appear in other contexts within statistical mechanics. [Fel70, Vol. I, pp. 39ff, 188ff]

The state space of the process is naturally identified with the set partitions of $[n]$ into m subsets with β_i elements in the i th subset. The state may also be associated with a Young tableau of fixed "shape" β , in which the contents $1, 2, \dots, n$, are decreasing in each row, but obey no order condition in the columns.

The process is a symmetric ergodic Markov chain. The distribution of the state therefore converges to the uniform distribution on the state space. How many steps are required in order that the true distribution of the process be close to the uniform stationary distribution? This question has been answered with tightest possible bounds by Diaconis and Shahshahani in the classical case when $m = 2$ [DS87]. (In an earlier result they also handled the process that results if we allow $n = m$ [DS81], which gives a random walk on the set of permutations of $[n]$ generated by applying random transpositions. Here we deal only with the model where $n > m$.)

The techniques used here provide elementary arguments matching the existing bounds when $m = 2$, and generalizing to $m > 2$. However they are weak for these larger m . We give a coupling for the general process, and give bounds on the time to stationarity in the two-urn and many-urn cases.

2.2.1 A Coupling for Bernoulli-Laplace Processes

Let Z be a Bernoulli-Laplace- (n, m, β) process. We describe a coupling (X, Y) for Z .

1. Let $\beta = (\beta_1, \beta_2, \dots, \beta_m)$. Start X with the first β_1 balls in the first urn, the next β_2 balls in the second urn, and so forth. Start Y in a state chosen uniformly from its possible states. Call a ball in either process coupled to its counterpart in the other process if both are in corresponding urns in the two processes.

2. Choose two urns $\{u, v\}$ uniformly at random from the $\binom{m}{2}$ such pairs. Choose the corresponding two urns in process Y .
3. Choose a ball a at random from urn u in process X . If ball a is coupled, choose the same ball $a' = a$ from urn u in process Y . If ball a is uncoupled, choose a ball a' uniformly from the uncoupled balls in urn u of process Y .
4. Choose a ball b at random from urn v in process X . If ball b is coupled, choose the same ball $b' = b$ from urn v in process Y . If ball b is uncoupled, choose a ball b' uniformly from the uncoupled balls in urn v of process Y .
5. Swap balls a and b in process X . Swap balls a' and b' in process Y .
6. Let T be the first time that all balls are coupled.

Theorem 2.9 *The procedure above gives a proper coupling (X, Y) for the (n, m, β) Bernoulli-Laplace process, for $n > m \geq 2$.*

Proof: It is clear that the process X is identical to Z , and that Y starts in the stationary distribution of Z . We must show that the transition matrix of Y is identical to Z , and that Y eventually meets X with probability 1, coinciding thereafter.

To see that Y moves with the same transition probabilities as Z , notice that (a) the urns are chosen with the same probabilities as for Z , and (b) when a ball is chosen from an urn, either it was coupled and the corresponding ball was chosen in X , or it was uncoupled and was chosen uniformly from amongst the uncoupled balls in the same urn in Y . Thus, a ball within a chosen urn i is picked with probability $\frac{1}{\beta_i}$ if it is coupled, while if it is not coupled it is chosen with probability $\frac{1}{k} \times \frac{k}{\beta_i} = \frac{1}{\beta_i}$, where k is the number of uncoupled balls in urn i . In both cases the ball is chosen with the correct probability.

To see that Y eventually meets X with probability 1, first notice that once a ball is coupled it remains coupled, and, as we show in the next paragraph, there is always a finite sequence of moves, having nonzero probability, that reduces the number of uncoupled balls. It follows that there is a finite sequence of moves, with nonzero probability, making X and Y coincide. Therefore they eventually meet with probability 1.

Assume for the moment that all β_i are at least 2. Let a be an uncoupled ball in some urn u of X . Then, by the definition of a "coupled" ball, a resides in some other urn, v , of Y . Since all β_i are at least 2, there is at least one other ball b' in urn v of Y . If b' is coupled, let b be its mate in urn v of X , otherwise let b be any uncoupled ball in urn v of X . Let a' be any uncoupled ball in urn u of Y . (There must be at least one, since there is one uncoupled ball, namely a , in urn u of X .) Now the move that swaps a and b in X , and a' and b' in Y , couples the previously uncoupled balls labelled a (and possibly others), while leaving coupled all previously coupled balls. A similar

sequence of two moves works even when only $\beta_1 \geq 2$ (guaranteed by $n > m$), by first getting a to urn 1 in Y , and then doing the move just described. \square

2.2.2 Tight Analysis of the Two-Urn Case

The time to stationarity in the two-urn case is easy to analyze. A “coupon collector” argument yields the following bound.

Theorem 2.10 *For the $(n, 2, \beta)$ Bernoulli-Laplace process, with $n > 2$, $\beta_1 \geq \beta_2$, we have*

$$\|\pi_k - \pi\| \leq e^{-c} \text{ whenever } k \geq \frac{\beta_1 \beta_2}{n-2} (\ln(\beta_1 \beta_2/n) + c).$$

Proof: In the coupling described before, let $E_{i,k}$ be the event that the i th ball is still uncoupled after the k th step. The event $\{T > k\}$ is precisely the union $\bigcup_i E_{i,k}$, where the union is taken over all of the balls i in the second urn of process X after the k th step. So $\Pr\{T > k\} \leq \sum_i \Pr\{E_{i,k}\} = \beta_2 \Pr\{E_{i_0,k}\}$, where i_0 is any one ball in the second urn. At time $k = 0$ we have $\Pr\{E_{i_0,0}\} = \beta_1/n$. Now consider that, in any given step, a given uncoupled ball (necessarily in urn 1 in one of the two processes) can get coupled in one of 2 ways: Either it stays fixed and its counterpart in the other process moves into the corresponding urn, which occurs with probability $\frac{\beta_1-1}{\beta_1} \times \frac{1}{\beta_2}$; or it moves into the other urn and its counterpart in the other process stays fixed, which occurs with probability $\frac{1}{\beta_1} \times \frac{\beta_2-1}{\beta_2}$. We thus find that a given uncoupled ball has probability exactly $p = \frac{(\beta_1-1) + (\beta_2-1)}{\beta_1 \beta_2} = \frac{n-2}{\beta_1 \beta_2}$ of being coupled at any step, and thus $\Pr\{E_{i_0,k}\} \leq \frac{\beta_1}{n} (1-p)^k$. So $\Pr\{T > k\} \leq \frac{\beta_1 \beta_2}{n} (1-p)^k$. Writing this in terms of the exponential, using the Taylor expansion for the logarithm, and applying the Coupling Inequality then yields the stated bound. \square

Remark: The case when $\beta_1 = \beta_2 = \frac{n}{2}$ gives the largest value of this bound on the time to stationarity. In this case, we have total variation distance not exceeding e^{-c} after $k > \frac{n^2}{4(n-2)} (\ln n - \ln 4 + c)$, or about $\frac{1}{4}n(\ln n + c)$ steps. This matches, within a factor of 2, the bound obtained by Diaconis and Shahshahani in [DS87], but our argument is elementary. Their bound is obtained using representation theory which we discuss briefly later. The properties they rely on no longer hold for the case of $n > m \geq 3$. We can give some analysis of the coupling, and we give some new geometric techniques based on the underlying graph. \square

2.2.3 Weak Analysis of the General Case

The preceding argument generalises easily to handle more than two urns each containing at least two balls. In this situation, each uncoupled ball has some chance of being coupled in the next step, as in the two-urn case.

Theorem 2.11 Let Z be an (n, m, β) Bernoulli-Laplace process, with $n > 2$, $\beta_1 \geq \beta_2 \geq \dots \geq \beta_m \geq 2$. Let π_k denote the distribution of the state of Z after the k th step, let π denote the stationary distribution, and let $\|\pi_k - \pi\|$ denote the total variation distance between them. Then we have

$$\|\pi_k - \pi\| \leq e^{-c}$$

whenever

$$k \geq \frac{\binom{m}{2}\beta_1\beta_2}{\beta_1 + \beta_2 - 2} (\ln(n - \beta_1) + c).$$

Proof: Again let $E_{i,k}$ be the event that the i th ball is still uncoupled after the k th step. Let I_k be the set of all balls in process Z , except those that lie in urn 1 immediately after the k th step of the coupling. Note that for all k , $|I_k| = n - \beta_1$. The event $\{T > k\}$ is precisely the union $\bigcup_{i \in I_k} E_{i,k}$, since if after step k all balls in the set I_k are coupled, then necessarily those in the first urn are coupled as well. So $\Pr\{T > k\} \leq \sum_{i \in I_k} \Pr\{E_{i,k}\} \leq (n - \beta_1)(1 - p)^k$, where p is a lower bound on the probability that a given uncoupled ball becomes coupled in a single step of the process. To obtain a value for p , notice that in any given step, a given uncoupled ball b can become coupled if the following sequence of events occurs.

1. A pair of urns $\{u, v\}$ where urn u contains the ball b in one process and urn v contains the counterpart ball b in the other process is chosen. This happens with probability $1/\binom{m}{2}$.
2. Then either the ball in u is chosen to move into the urn v (probability $1/\beta_u$), while the counterpart remains fixed (probability $(\beta_v - 1)/\beta_v$), or alternatively, the ball in u stays fixed (probability $(\beta_u - 1)/\beta_u$), and the counterpart in v moves to u (probability $1/\beta_v$).

Thus an uncoupled ball has probability at least

$$p = \min_{u,v} (\beta_u + \beta_v - 2) / \left(\binom{m}{2} \beta_u \beta_v \right) = (\beta_1 + \beta_2 - 2) / \left(\binom{m}{2} \beta_1 \beta_2 \right)$$

of being coupled at any step. Using this value of p with $\Pr\{T > k\} \leq (n - \beta_1)(1 - p)^k$, yields the stated bound. \blacksquare

Example 2.12 For the case of n balls distributed evenly in $m = 3$ urns, the theorem says that approximately $\frac{1}{4}n \ln n$ swaps are sufficient to get near uniformity. More generally, for n balls distributed evenly in any fixed number m of urns, $O(n \ln n)$ steps are sufficient to get within any fixed variation from uniformity. But if m is allowed to increase, this bound grows as m^2 . This is wrong, but seems inherent in this line of analysis. \square

A similar, but even weaker, argument can be used to give bounds when one of the urns contains a single ball. When more than one urn contains only one ball, a slightly different argument is needed. For in this case, uncoupled balls that reside in one of the single-ball urns and whose counterparts

in the other process are also in single-ball urns have no chance of getting coupled in a single step. They do have a chance, however, of getting coupled in two steps by first moving (or having their counterparts move) to an urn with many balls, and then coupling in one more step. Letting ℓ denote the number of urns containing at least two balls, the probability of this two-step event can be bounded below by $\frac{2\ell}{\binom{m}{2}}p$, where p is that defined in the proof above. This places a multiplicative factor of $\frac{(\ell)}{\ell}$ on the time to stationarity obtained above. Clearly this analysis is not very tight, since putting $\ell = m$ does not return us the bound of Theorem 2.11, but instead sacrifices a large factor of about $\frac{m}{2}$. A significantly tighter analysis using this coupling is not evident.

2.3 Markov Chains based on Groups

In the remainder of this chapter we explore random walks based on groups and group actions, which is the natural setting in which to work more generally with urn models. We will use the Bernoulli-Laplace model as a running example. Before proceeding, we need to introduce a substantial amount of terminology.

2.3.1 Transitive Group Actions

We assume the reader is familiar with the basic properties of a group. Otherwise the reader may refer to Hungerford's book [Hun74]. A more elementary text will do. Here when we speak of a group, we will mean a finite group.

Let X be a finite set. A (left) action of a group Γ on X is a function α mapping $\Gamma \times X$ to X such that for all $z \in X$, $g_1, g_2 \in \Gamma$

$$\alpha(\text{id}, z) = z \quad \text{and} \quad \alpha(g_2, \alpha(g_1, z)) = \alpha(g_2g_1, z),$$

where id denotes the identity element. It is natural to omit α and write simply: gz for $g \in \Gamma$, $z \in X$. This notation is ambiguous; the same group may have several different actions on a set. But in context, it poses no problems.

Given an action, and elements $z, y \in X$, say that z may be mapped to y , if $gz = y$ for some $g \in \Gamma$. This is an equivalence relation on X , and the equivalence classes are called the orbits of the action. The equivalence class containing z is called the orbit of z . If all elements may be mapped to each other, that is, if the action has only one orbit, the action is called transitive. *From this point on, we restrict our attention to transitive actions.*

Example 2.13 A group Γ acts transitively on itself through its multiplication map: Let $X = \Gamma$ and let $\alpha(g, z)$ be the group element gz obtained by multiplying g and z . This action is called left translation, and the corresponding right action of Γ is called right translation. \square

Example 2.14 We get a transitive action of the symmetric group S_n on $[n]$ by defining for $\sigma \in S_n$, $z \in [n]$, $\sigma z = \sigma(z)$, the image of the element z under the permutation σ . \square

Example 2.15 The symmetric group S_n acts "elementwise" on the set of all k -element subsets of $[n]$. For a k -set $z = \{z_1, z_2, \dots, z_k\}$ and $\sigma \in S_n$, define

$$\sigma z = \{\sigma(z_1), \sigma(z_2), \dots, \sigma(z_k)\},$$

the set of the images of the elements of z under the permutation σ . This gives a transitive action. \square

Remark: If Γ acts on X consider the function $g(z) = gz$. This is one-to-one, for if $gz_1 = gz_2$ then $z_1 = z_2$; it is also surjection since $g(g^{-1}y) = y$. This allows one to view a group action as a homomorphism of Γ into the group of permutations of X . \square

If x and y are any two elements of X , let Γ_{xy} denote the (non-empty) set of elements mapping x to y . That is, $\Gamma_{xy} = \{g \mid gx = y\}$.

A group element g such that $gx = x$ is said to fix x . The set of all elements that fix x forms a subgroup of Γ called the stabilizer or isotropy subgroup of x and denoted $\text{Stab}(x)$. Note that $\text{Stab}(x) = \Gamma_{xx}$.

Lemma 2.16 Let Γ act transitively on X . For given x, y , let $g_0 \in \Gamma_{xy}$ be some designated element mapping x to y . Then every element in Γ_{xy} may be written uniquely as g_0n for some $n \in \text{Stab}(x)$.

Proof: Define the function $f(n) = g_0n$. We wish to show this gives a bijection of $\text{Stab}(x)$ onto Γ_{xy} . It is clearly a one-one function by cancellation. Moreover, the image of $\text{Stab}(x)$ is clearly contained in Γ_{xy} .

It remains only to show that every element in Γ_{xy} is the image $f(n)$ of some element in $n \in \text{Stab}(x)$. Suppose $g \in \Gamma_{xy}$ so that $gx = y$. Then

$$(g_0^{-1}g)x = g_0^{-1}(gx) = g_0^{-1}y = g_0^{-1}(g_0x) = (g_0^{-1}g_0)x = x.$$

So $n = g_0^{-1}g$ is in $\text{Stab}(x)$, and $f(n) = g_0n = g$. \square

Corollary 2.17 It follows that for every $x, y \in X$

$$|\Gamma_{xy}| = |\text{Stab}(x)| = \frac{|\Gamma|}{|X|}.$$

Note that the latter is independent of x and y .

Proof: The first equality is immediate from Lemma 2.16. For the second, fix an $z \in X$. Clearly $\Gamma = \bigcup_{y \in X} \Gamma_{zy}$ (since every g maps z to some y), and the Γ_{zy} are disjoint (since g cannot map z to two different y 's). From the first equality, each Γ_{zy} has size $|\text{Stab}(z)|$, so $|\Gamma| = |X||\text{Stab}(z)|$. \blacksquare

Remark: The latter equality is a special case of a more general theorem.

$$|\text{Stab}(z)| = \frac{|\Gamma|}{|\text{Orb}(z)|},$$

where $\text{Orb}(z)$ denotes the orbit of z . For a proof, see [Hun74, Theorem 4.3]. \square

Here is another useful way of viewing things that we shall use. Designate an element $z_0 \in X$ and let $N = \text{Stab}(z_0)$ be the isotropy subgroup. For $y \in X$, the sets Γ_{z_0y} are precisely the cosets of $\text{Stab}(z_0)$ in Γ . If g maps z_0 to y then every element of gN maps z to y . If g and h both map z_0 to y then $hN = gN$.

If we designate an element g_y in each distinct coset Γ_{z_0y} , we get a system of coset representatives. This gives a mapping of X to the coset space Γ/N that respects the action of Γ : the coset corresponding to z_0 is N itself, and in general for $y \in X$ the corresponding coset is $g_yN = \Gamma_{z_0y}$. Every element of Γ can be written uniquely as $g_y n$ for some y and some $n \in N$, and for any z, y , every element of Γ_{zy} can be written $g_y n g_z^{-1}$ for some $n \in N$.

All of the isotropy subgroups are conjugate since $\text{Stab}(y) = g_y^{-1} \text{Stab}(z_0) g_y$. This means our choice of z_0 was not important, since this conjugacy gives an isomorphism of the group.

2.3.2 Cayley Graphs

Let Γ act transitively on X . Let S be a symmetric set of generators for Γ : S generates Γ and $s \in S$ implies $s^{-1} \in S$. The simple graph $G = (X, E)$ whose vertex set is X , and whose edges are $E = \{\{z, y\} \mid z, y \in X, z \neq y, y = sz \text{ for some } s \in S\}$ is the (Cayley) graph of the group action with respect to the generators S , and we denote it $\text{Cayley}(\Gamma, X, S)$, where the action is understood from context. This graph is connected since the action is transitive. In the special case when $X = \Gamma$, the group itself, and the action is given by translation, the graph $\text{Cayley}(\Gamma, \Gamma, S)$ is known simply as the Cayley graph of the group (with respect to the generators S).

Example 2.18 Let Γ be S_n . Let S be the set of all transpositions. Then $\text{Cayley}(\Gamma, \Gamma, S)$ is the graph with a vertex for each arrangement of $[n] = \{1, 2, \dots, n\}$ with an edge between two vertices if their arrangements differ by one swap. Fix a k and let X be the set of k -sets of $[n]$. Let Γ act elementwise on the k -sets. Then $\text{Cayley}(\Gamma, X, S)$ is the graph with a vertex for each of the $\binom{n}{k}$ k -sets of $[n]$ and an edge between two vertices if the symmetric difference of the corresponding k -sets has cardinality 2. \square

2.3.3 Vertex-Transitive Graphs

If $G = (V, E)$ is an undirected simple graph, its automorphism group, denoted $\text{Aut}(G)$, is the set of all permutations g of the vertices such that $g(E) = E$, where $g(E)$ denotes the set composed of the elements $\{g(v), g(w)\}$ for each $\{v, w\} \in E$.

The graph G is called vertex-transitive if $\text{Aut}(G)$ acts transitively on the vertices, or in other words, if for every pair of vertices $v, w \in V$ there is an automorphism $g \in \text{Aut}(G)$ such that $g(v) = w$.

Intuitively, this means that the graph "looks the same" from every vertex.

It is well known that Cayley graphs of groups are always vertex-transitive [Big74]. We include a proof for completeness.

Theorem 2.19 *If Γ is any group and S is any symmetric set of generators of Γ , then $G = \text{Cayley}(\Gamma, \Gamma, S)$ is vertex-transitive.*

Proof: Γ acts transitively on itself by left translation $z \mapsto gz$ for $z \in \Gamma$, and it is easy to see that each $g \in \Gamma$ gives an automorphism of the graph by this action. Let $G = (\Gamma, E) = \text{Cayley}(\Gamma, \Gamma, S)$. Then

$$\begin{aligned} (z, y) \in E &\Leftrightarrow yz^{-1} \in S \\ &\Leftrightarrow ygg^{-1}z^{-1} \in S \\ &\Leftrightarrow (yg)(zg)^{-1} \in S \\ &\Leftrightarrow ((gz), (gy)) \in E \end{aligned}$$

It follows that the permutations of Γ given by $z \mapsto gz$ are all in $\text{Aut}(G)$, and since Γ acts transitively so does $\text{Aut}(G)$. \square

Example 2.20 The same does not hold generally for Cayley graphs of group actions. Cayley graphs of group actions are not always vertex-transitive. Take $\Gamma = S_4$, and let X be the set of 2-element subsets of $[4]$. For generators take $S = \{(12), (234), (432)\}$. The graph has $\binom{4}{2} = 6$ vertices so that it can easily be drawn. One can see by inspection that $\text{Cayley}(\Gamma, X, S)$ has no automorphism mapping the vertex for the set $\{1, 2\}$ which has degree 2, to the set $\{1, 3\}$, which has degree 3.

It is possible to define Cayley graphs of group actions as graphs with self-loops and multiple edges so that they are, at least, regular. One can check that the preceding example is still not vertex-transitive under that definition. \square

Cayley graphs of group actions are vertex-transitive only in certain cases. For example, if Γ acts transitively on X and the isotropy subgroup N is a normal subgroup of Γ , then $\text{Cayley}(\Gamma, X, S)$ is isomorphic to $\text{Cayley}(\Gamma/N, \Gamma/N, S')$, the graph of the quotient group Γ/N under the generators $S' = \bigcup_{s \in S} sN$. A more interesting situation is the following.

Theorem 2.21 Let Γ act transitively on X , and let S be a symmetric set of generators for Γ . If S is closed under conjugacy then $G = \text{Cayley}(\Gamma, X, S)$ is vertex-transitive.

Proof: Let $z \mapsto gz$ be the action of Γ on X . We know this is a transitive action. We show that every $g \in \Gamma$ gives automorphism of $G = (X, E)$ through this action.

$$\begin{aligned}
 (z, y) \in E &\Leftrightarrow y \in Sz \text{ and } z \neq y \\
 &\Leftrightarrow y \in g^{-1}Sgz \text{ and } z \neq y \\
 &\Leftrightarrow gy \in Sgz \text{ and } z \neq y \\
 &\Leftrightarrow (gz, gy) \in E \text{ and } z \neq y \\
 &\Leftrightarrow (gz, gy) \in E \text{ and } gz \neq gy.
 \end{aligned}$$

■

Example 2.22 In S_n , the transpositions form a single conjugacy class. Thus the graph of k -sets described in Example 2.18 is vertex-transitive. \square

2.3.4 Chains Based on Groups

A Markov chain is a random walk on a group Γ if its transition matrix P has entries given by

$$P_{x,y} = p(yx^{-1}) \quad (\forall x, y \in \Gamma)$$

where p is some probability distribution on the group Γ . The distribution p is called the transition measure.

The chain P is ergodic, provided that the support of the transition measure $\text{Supp}(p) = \{g \mid p(g) > 0\}$ generates Γ . Moreover, it is easy to see that the transition matrix P of any random walk on a group is doubly stochastic. This means that its stationary distribution is the uniform distribution on Γ .

Because the stationary distribution is uniform, P will be time-reversible precisely if P is symmetric. This can be stated in terms of the transition measure: P is time-reversible, and in fact symmetric, iff $p(g) = p(g^{-1})$. In this case, we will call P a symmetric random walk on Γ .

If the group Γ acts transitively on a finite set X , any random walk P on Γ induces a Markov chain P' on X as follows. Fix an element $z_0 \in X$. When the walk P is in state $g \in \Gamma$, we specify that the Markov chain P' is in state $gz_0 \in X$. It is easy to see that if z_k is the state of the chain P' at time k , then the next state is determined by $z_{k+1} = gz_k$ where $g \in \Gamma$ is chosen according to the transition measure p of the original walk P . We say that P' is P induced directly to X . The entries of the transition matrix P' are:

$$P'_{x,y} = p(\Gamma_{xy}) = p(g_y N g_x^{-1}) \quad (\forall x, y \in X), \quad (2.2)$$

where $N = \text{Stab}(x_0)$. The underlying graph of any symmetric random walk on Γ induced directly to X is the Cayley graph of the group's action. We can and will use this to our advantage in bounding eigenvalues.

For a certain class of transition measures, there is another way to induce a walk on X through a group's action. This alternate method is less direct but yields a chain with greater symmetry. A measure p (or more generally any function on Γ) is called N -bi-invariant if for every $g \in \Gamma$ and for all $n_1, n_2 \in N$,

$$p(n_1 g n_2) = p(g).$$

Again, fix an $x_0 \in X$ and let $N = \text{Stab}(x_0)$. If P is a random walk on Γ given by an N -bi-invariant transition measure p , then we can define a Markov chain \tilde{P} on X by

$$\tilde{P}_{x,y} \stackrel{\text{def}}{=} p(g_x^{-1} g_y N). \quad (2.3)$$

This is well-defined, and does not depend on our choice of coset representatives since any elements of the cosets $g_x N$ and $g_y N$ can be represented as $g_x n_1 g_y n_2$ for $n_1, n_2 \in N$, whence for some $n_3 \in N$ we have

$$p((g_x n_1)^{-1} (g_y n_2) N) = p(n_1 g_x^{-1} g_y N n_3) = p(g_x^{-1} g_y N),$$

by the bi-invariance of p .

We say that \tilde{P} is P induced symmetrically to X . The chain has the following intuitive description. At each step, we imagine we are at x_0 , and we first choose a new position as if we were at x_0 i.e., we choose a position $g x_0$ choosing $g \in \Gamma$ according to $p(g)$. Then, using the action of g_x , we translate this whole move $(x_0, g x_0)$ into a move from x to $y = g_x g x_0$. It is easy to see that to move from x to a given element y it is necessary and sufficient that $g \in g_x^{-1} g_y N$.

The chain satisfies the following symmetry condition.

$$\tilde{P}_{g_x x, g_y y} = \tilde{P}_{x, y} \quad \forall g \in \Gamma. \quad (2.4)$$

and it is easily verified that any Markov chain $P_{x,y}$ on $X = \Gamma/N$ that satisfies this condition is generated by an N -bi-invariant transition measure given by $p(g) = \frac{1}{|N|} P_{x_0, g x_0}$.

Theorem 2.23 *Let Γ act transitively on X with isotropy subgroup $N = \text{Stab}(x_0)$. Let p be a transition measure on Γ that is constant on conjugacy classes. The chain P' induced directly to X is the same as the chain \tilde{Q} induced symmetrically to X starting from the N -bi-invariant transition measure*

$$q(g) = \frac{1}{|N|} p(gN).$$

Proof: First, let us see that \tilde{Q} is well defined by showing that $q(g)$ is indeed N -bi-invariant. Right-invariance is clear. On the left, for $n \in N$ we have

$$q(ng) = \frac{1}{|N|} p(ngN) = \frac{1}{|N|} p(ngNn^{-1}) = \frac{1}{|N|} p(gN) = q(g),$$

since p is constant under conjugacy and N is a subgroup.

Now we want to verify that $\tilde{Q}_{x,y} = P'_{x,y}$, for all $x,y \in X$. To do this, we first show that the chain P' satisfies the symmetry condition (2.4). Then, by that symmetry, it will suffice to consider only x_0 and show that $P_{x_0,y} = \tilde{Q}_{x_0,y}$, for all $y \in X$.

The chain P' satisfies the symmetry condition because p is constant on conjugacy classes. Specifically, we have

$$P'_{gx,gy} = p((gg_y)N(gg_x)^{-1}) = p(gg_yN g_x^{-1}g^{-1}) = p(g_yN g_x^{-1}) = P'_{x,y}$$

for every $g \in \Gamma$.

Now, restricting our attention to x_0 , we have

$$P_{x_0,y} = p(g_yN g_{x_0}) = p(g_yN) = q(g_yN) = q(g_{x_0}^{-1}g_yN) = \tilde{Q}_{x_0,y}.$$

Here in the second equality, we have used the fact that g_{x_0} is in N . In the fourth equality, the same fact was combined with the N -invariance of q . \blacksquare

Example 2.24 Let P be the random walk on $\Gamma = S_n$, generated using $p(\text{id}) = 1/n$, $p(g) = 2/n^2$ if g is a transposition and $p(g) = 0$ otherwise. Let X be the k -sets of n , and let $x_0 = \{1, 2, \dots, k\}$. The group Γ acts transitively on X elementwise; the isotropy subgroup $N = \text{Stab}(x_0)$ is isomorphic to $S_k \times S_{n-k}$. The directly induced walk is P' , whose nonzero entries are $P'_{x,y} = 2/n^2$ if the symmetric difference of x and y has cardinality 2, and $P'_{x,x} = 1 - (2k(n-k)/n^2)$ for each x . The underlying graph is $\text{Cayley}(\Gamma, X, S)$, where S is the set of transpositions.

The theorem says this is the same as the chain \tilde{Q} induced symmetrically from the N -bi-invariant transition measure $q(g) = \frac{1}{|N|}p(gN)$. For example, $Q_{x,x} = q(N) = (\frac{1}{n} + \frac{2}{n^2}(\frac{k}{2}) + \frac{2}{n^2}(\frac{n-k}{2})) = 1 - (2k(n-k)/n^2)$, as required.

Note that this chain is not quite the same as the Bernoulli-Laplace model; it differs in the holding probabilities. The Bernoulli-Laplace model corresponds to the walk induced symmetrically from the N -bi-invariant measure determined by $q(gN) = \frac{1}{k(n-k)}$ if gN is a coset corresponding to a set within a single exchange of x_0 . \square

Remark: This gives a partial answer to a question posed by Diaconis and Shahshahani [DS87, p. 213]: when can the directly induced walk from a given transition measure p be viewed as the symmetrically induced walk of some N -bi-invariant measure q ? They give techniques for dealing with the latter in some cases (which we discuss briefly in the next section). This gives a class of cases in which the same techniques will apply to the directly induced walk. \square

The following should be noted. A coupling for a random walk on Γ induces couplings for the induced walks in the obvious way (both in the direct and symmetric case). When the coupling for Γ succeeds, so does the induced coupling. So a maximal coupling for Γ induces a coupling for the

induced walk that converges at least as fast. This gives us the following simple theorem. (The theorem also has other simple proofs in terms of properties of variation distance.)

Theorem 2.25 *Both the direct and symmetrically induced walks converge in variation distance at least as fast as the random walk on the group.*

2.3.5 On the Harmonic Analysis Approach

Diaconis [Dia88] has applied representation theory to give good bounds on the time to stationarity for a number of random walks on groups. The advantage of such results is that they typically give bounds on the full spectrum, and lead to sharp convergence rate bounds; in many cases matching lower bounds on the convergence rate can be found. Similar techniques have been used by Flatto, Odlyzko, and Wales [FOW85] to analyze the time to hit a particular state, and by Matthews [Mat85] to analyze the time to hit all states (covering time).

The idea behind the harmonic analysis approach can be summarized as follows. For background, the reader should consult [Dia88]. If P is a random walk on a group, the k th-step distribution $\pi_k = \pi_0 P^k$ may also be expressed as the k -fold convolution of the transition measure. That is to say

$$\pi_k(y) = p^{-k}(y) = \sum_{z \in \Gamma} p(yz^{-1}) p^{*(k-1)}(z).$$

At any representation of the group, the Fourier transform has the property that the transform of a convolution of two functions is the product of the individual transforms. This converts the complicated convolution operation on the group to a matrix multiplication. Using a few basic properties of representations, one gets the following lemma.

Lemma 2.26 (Upper Bound Lemma for Groups and Actions) [DS81, Dia88] *If P is the random walk on Γ with transition measure p then the k -th step distribution π_k of P satisfies*

$$\|\pi_k - U\|^2 \leq \frac{1}{4} \sum_{\rho} \deg \rho \text{Tr}[(\hat{p}(\rho))^k (\hat{p}(\rho)^*)^k].$$

where $\hat{p}(\rho)$ represents the Fourier transform at the representation ρ and the sum is over all non-trivial irreducible representations ρ . (Here $*$ indicates the conjugate-transpose operation.)

If N is a subgroup of Γ and the transition measure p is N bi-invariant, the k th-step distribution $\widetilde{\pi}_k$ of the symmetrically induced walk \tilde{P} on $X = \Gamma/N$ satisfies

$$\|\widetilde{\pi}_k - U\|^2 \leq \frac{1}{4} \sum_{\rho} \deg \rho \text{Tr}[(\hat{p}(\rho))^k (\hat{p}(\rho)^*)^k].$$

where the sum is over all non-trivial irreducible representations ρ that occur in $L(X)$, the representation given by the set of all complex-valued functions on X .

This is analogous to the bound of Theorem 1.9. The eigenvalues of $\hat{p}(\rho)$ are those of the Markov transition matrix P appearing with multiplicity given by the degree of ρ . For details see Diaconis [Dia88, pp. 48-49]. Essentially the same relationship is discussed by Babai [Bab79] with respect to the eigenvalues of the Cayley graph's adjacency matrix.

The repeated multiplication of the Fourier transform matrices $\hat{p}(\rho)$ is not necessarily any more analyzable than the multiplication of the transition matrix of the Markov chain. But under certain circumstances one can insure that the matrices $\hat{p}(\rho)$ that arise have special structure that makes the analysis simpler. Except for a few special cases (e.g. [BD89]), here is an essentially complete list of the circumstances that Diaconis and others have been able to take advantage of.

- **Abelian groups.** When Γ is abelian, all of its irreducible representations are 1-dimensional. That is, the Fourier transform matrices $\hat{p}(\rho)$ involved are, in fact, (complex) scalars.
- **Transition measures constant on conjugacy classes.** When the transition measure p is constant on conjugacy classes the Fourier transforms at every irreducible representation are almost-scalar; they are of the form cI where I is the identity matrix of order equal to the degree of the representation. On specific groups in which the character theory is well-developed, this can make the problem tractable.
- **Symmetrically induced walks when (Γ, N) forms a Gelfand Pair.** The group Γ and subgroup N form a Gelfand pair when the convolution of N -bi-invariant functions on Γ is commutative. Under these circumstances Diaconis and Shahshahani [DS87, Dia88] give a beautiful treatment for walks \tilde{P} on Γ/N induced symmetrically from an N -bi-invariant transition measure p on Γ . They take advantage of the fact that, in this case, the Fourier transform of p at any irreducible representation is, in some basis, a matrix with a single nonzero entry in position $(1, 1)$. Again, where the representation theory allows calculation of, or at least bounds on, these $(1, 1)$ entries, there is hope of getting good convergence bounds. By Theorem 2.23, a similar analysis applies to directly induced walks on Γ/N when (Γ, N) is a Gelfand pair.

When these techniques apply, and a tight elementary argument is not apparent, they should be used. They seem to give tight bounds for many cases. In addition there may be other reasons that one wants to do the harmonic analysis, since there are other applications of the representation theory to the analysis of certain statistical data that can be viewed as data on groups. [Dia88]

The new methods given in the remaining sections can be applied when the harmonic analysis does not prove tractable.

2.3.6 Magnification Bounds

In this section we present some new bounds on the magnification of some classes of underlying graphs that arise from random walks on groups. These bounds use very little except basic group structure.

yielding results that are in terms of simpler quantities, "diameters." Such bounds are not usually tight, but are simple to calculate and give reasonable results.

The basic structural properties of group actions allow us to extend a theorem of Aldous [Ald87] to obtain the following generalization. The proof closely matches Aldous's proof. Recall the definition of magnification from Section 1.5.2.

Theorem 2.27 (Magnification in Cayley Graphs) Suppose Γ acts transitively on X , and let $G = \text{Cayley}(\Gamma, X, S)$. Let Δ any number such that every element of the group Γ can be written with at most Δ generators from S . Then G has magnification

$$c \geq \frac{1}{2\Delta}.$$

Proof: Let $A \subset X$ with $|A| \leq |X|/2$. Let $\delta(a, b)$ be the function that is 1 if $a = b$, 0 otherwise. By Corollary 2.17

$$|\Gamma_{xy}| = |\text{Stab}(x)| = \frac{|\Gamma|}{|X|}.$$

Therefore we have,

$$\begin{aligned} \sum_{g \in \Gamma} |gA \cap A| &= \sum_{g \in \Gamma} \sum_{x \in A} \sum_{y \in A} \delta(gx, y) \\ &= \sum_{x \in A} \sum_{y \in A} \sum_{g \in \Gamma} \delta(gx, y) \\ &= \sum_{x \in A} \sum_{y \in A} |\Gamma_{xy}| \\ &= \frac{|A|^2 |\Gamma|}{|X|}. \end{aligned}$$

Thus for some \hat{g} in Γ , we must have $|\hat{g}A \cap A| \leq \frac{|A|^2}{|X|}$ (the average), which implies that $|\hat{g}A - A| \geq |A|/2$.

Now write $\hat{g} = h_d h_{d-1} \cdots h_2 h_1$ for some $d \leq \Delta$ with each $h_i \in S$, and let $g_0 = \text{id}$, $g_i = h_i h_{i-1} \cdots h_2 h_1$ for $1 \leq i \leq d$. Then we have

$$|\hat{g}A - A| \leq \sum_{i=1}^d |g_i A - g_{i-1} A| = \sum_{i=1}^d |h_i A - A|.$$

So there must be an element \hat{h} among the h_i such that

$$|\hat{h}A - A| \geq \frac{|\hat{g}A - A|}{\Delta} \geq \frac{|A|}{2\Delta}.$$

This, in turn, gives $c \geq \frac{1}{2\Delta}$. \blacksquare

When the set X is taken to be Γ itself, and the action is given by the usual multiplication in Γ , the graph G is simply the Cayley graph of the group under S , and Δ is the diameter of the graph G .

This is a special case for which Aldous proved the same eigenvalue bound. When X is not Γ , the quantity Δ in the lower bound is not necessarily the diameter of the graph.

We have seen that Cayley graphs of groups are vertex transitive. Babai [Bab90] has recently proved the following theorem which gives a generalization of Aldous's result in this direction. In this the quantity Δ corresponds to the diameter. We adapt his proof only slightly here. Note that Cayley graphs of group actions are not always vertex-transitive, so this does not completely supersede the previous theorem.

Theorem 2.28 (Magnification in Vertex-Transitive Graphs) (From [Bab90]) *Let G be a vertex-transitive graph with diameter Δ . Then G has magnification*

$$c \geq 1/2\Delta.$$

Proof: Let $G = (V, E)$ and let $\Gamma = \text{Aut}(G)$ be the automorphism group of G . We know this acts transitively on V . Fix a vertex $v_0 \in V$ and let $N = \text{Stab}(v_0)$ be the isotropy subgroup. And for each element $v \in V$ let g_v be a representative element in the coset Γ_{v_0v} .

We construct a certain Cayley graph of the group Γ and show that the earlier theorem applied to this Cayley graph implies the result for G .

Let C be the graph with vertex set Γ in which we join two vertices $g, h \in \Gamma$ if $gv_0 = hv_0$ or if (gv_0, hv_0) is an edge in G .

The condition $gv_0 = hv_0$ is equivalent to the condition $g^{-1}h \in N$. Because g is an automorphism of G , (gv_0, hv_0) is an edge of G iff $(v_0, g^{-1}hv_0)$ is an edge of G . The latter is equivalent to the condition $g^{-1}h \in g_w N = \Gamma_{v_0w}$ for some w adjacent to v_0 . Using this one can see that $C = \text{Cayley}(\Gamma, \Gamma, S)$ where

$$S = \bigcup_{w=v_0 \text{ or } \{w, v_0\} \in E} g_w N.$$

The diameter of C is at most that of G , because if v_0, v_1, \dots, v_k is any path between v_0 and v_k in G , there is a corresponding path of the same length in C among the coset representatives $g_{v_0}, g_{v_1}, g_{v_2}, \dots, g_{v_k}$. (By vertex-transitivity it suffices to consider only paths from v_0 .) Hence the diameter of C is at most Δ . (It is not hard to see that the diameters are in fact equal.)

Let $f(g) : \Gamma \rightarrow V$ be the projection that takes g to gv_0 . We know by Lemma 2.16 that f maps the same number $|N|$ elements to any vertex v . Also, It is easy to see that this projection preserves neighborhoods in the sense that, $\text{Nbd}(f^{-1}(A)) = f^{-1}(\text{Nbd}(A))$, for $A \subset V$. It follows from Theorem 2.27 that

$$\frac{|\text{Nbd}(A)|}{|A|} = \frac{|f^{-1}(\text{Nbd}(A))|}{|f^{-1}(A)|} \geq 1/2\Delta.$$

■

Remark: Theorem 2.27 was obtained and presented in lectures in 1987, but is appearing in print for the first time. Babai's Theorem 2.28 supersedes a result in Diaconis [DS89] for distance-transitive

graphs; all such graphs are necessarily also vertex-transitive. See [Big74] for definitions. Babai's recent result was incorporated in the final draft. \square

2.3.7 Eigenvalue Bounds for the Chains

Based on Theorem 1.17 we can now give eigenvalue bounds for symmetric walks on groups, based on the magnification bounds given in the preceding section for their underlying graphs. If better bounds are known for the magnification of the underlying graph, one should appeal directly to Theorem 1.17.

For the following two theorems, we assume that Γ is a finite group, $p(\cdot)$ is a symmetric transition measure ($p(g) = p(g^{-1})$) whose support $S = \text{Supp}(p(\cdot))$ generates Γ , P is the associated walk on Γ , and Γ acts transitively on the set X with isotropy subgroup N .

Theorem 2.29 Let Δ be the diameter of $\text{Cayley}(\Gamma, \Gamma, S)$. The second largest eigenvalue of P satisfies

$$\lambda_1(P) \leq 1 - \frac{pc^2}{4 + 2c^2},$$

where $c = 1/2\Delta$ and $p = \min_{g \in S - \{\text{id}\}} p(g)$.

Let Δ' be the diameter of $\text{Cayley}(\Gamma, X, S)$ if S is closed under conjugacy, and let $\Delta' = \Delta$ otherwise. If P' is the directly induced walk on X , then $\lambda_1(P')$ satisfies

$$\lambda_1(P') \leq 1 - \frac{p'c^2}{4 + 2c^2},$$

with $p' = \min_{g \in S - N} p(gN)$ and $c = 1/2\Delta'$.

Proof: For the chain P , the underlying graph is $\text{Cayley}(\Gamma, \Gamma, S)$. For P' the underlying graph is $\text{Cayley}(\Gamma, X, S)$. In each case the value p is the minimum non-holding transition probability. The eigenvalue bounds are obtained by applying Theorem 1.17 together with Theorem 2.27.

We are justified in using the smaller value of Δ' when $\text{Supp}(p(\cdot))$ is closed under conjugacy, because Theorem 2.21 insures that the underlying graph is vertex-transitive. Thus we can apply Theorem 1.17 with Theorem 2.28 instead of Theorem 2.27. \square

Theorem 2.30 Let p be a symmetric N -bi-invariant measure with support S generating Γ , and let \tilde{P} be the symmetrically induced Markov chain on X . If Δ is the diameter of the underlying graph of \tilde{P} then

$$\lambda_1(\tilde{P}) \leq 1 - \frac{pc^2}{4 + 2c^2},$$

where $c = 1/2\Delta$, and $p = \min_{g \in S - N} p(gN)$.

Proof: The underlying graph is not necessarily the Cayley graph of the group action, but by the symmetry condition (2.4), it is necessarily vertex-transitive. Again, p is the minimum non-holding transition probability and the eigenvalue bounds are obtained by applying Theorem 1.17 with Theorem 2.28. \blacksquare

These bounds can be applied directly with Theorem 1.8 to get convergence bounds. We give some examples in the next section.

Remark: A similar bound holds for any symmetric ergodic Markov chain whose underlying graph is vertex-transitive by combining Theorem 2.28 and Theorem 1.17. But it seems that the most interesting cases are covered here. \square

2.3.8 Examples

We now show some examples and compare to tighter bounds where they are known. The convergence rates here are all obtained by converting the chain P in question to a strongly aperiodic chain $(I + P)/2$, and then applying the bounds of the previous section with Theorem 1.8. Note that converting to the strongly aperiodic chain decreases the minimum transition probability by a multiplicative factor of $\frac{1}{2}$, slowing the convergence bound by approximately a factor of 2.

The examples seem to display that the bounds are easily applicable, but usually not very tight.

Example 2.31 Simple Random Walk on the Affine Group. Consider the random walk on the affine group A_p when p is an odd prime and 2 is a primitive root modulo p . A_p has $p(p-1)$ elements and the 9 elements $\{(a, b) \mid a \in \{1, 2, p+1/2\}, b \in \{0, 1, -1\}\}$ are generators with $\Delta = O(p)$. We get an immediate $O(p^2 \ln p)$ convergence guarantee (for fixed positive $\epsilon < 1$). Diaconis [Dia88][Example 4, pp. 34-35] uses a neat trick to make the harmonic analysis work, and gets a slightly weaker bound that is $O(c(p)p^2 \ln p)$ steps, where $c(p)$ is any function that grows to infinity with p . By considering the projected walk obtained by tracking only the first coordinate a of each successive position (a, b) in the walk, it is not hard to see that cp^2 steps, where c is fixed, is not sufficient. Hildebrand [Hil90] has recently shown that $O(c(p)p^2)$ steps are sufficient, where again $c(p)$ is any function growing to infinity with p . \square

Example 2.32 Random Transpositions. Consider the random walk on the symmetric group generated by the measure $p(\text{id}) = 1/n$ and $p(g) = 2/n^2$ if g is a transposition. Here we have $\Delta = n - 1$, $p = \frac{2}{n^2}$, and an immediate eigenvalue bound of approximately $1 - \frac{1}{16n^4 + 2n^3}$. In actuality, $\lambda_1 = 1 - \frac{2}{n}$. The transpositions form a single conjugacy class, so the transition measure is constant on conjugacy classes. Using harmonic analysis and a full-spectrum bound, one finds $O(n \ln n)$ steps are sufficient to get a nearly uniform permutation, and this is tight [DS81] [Dia88]. Our second-eigenvalue bound gives a much weaker $O(n^3 \ln n)$ bound, but is a one line calculation. \square

Example 2.33 Transposition and $(n - 1)$ -cycle, Adjacent Transpositions. The symmetric group S_n is generated by a single transposition $(1 2)$ and the $(n - 1)$ -cycle $(2 3 4 \cdots n)$. Let p be the uniform distribution on $(1 2)$, $(2 3 4 \cdots n)$ and its inverse $(n (n - 1) \cdots 3 2)$. The diameter satisfies $\Delta = O(n^2)$, and for the strongly aperiodic form we have $p = 1/6$. The resulting bounds are $\mu_1 = \Omega(1/n^4)$, and convergence in $O(n^5 \ln n)$ steps. This is probably way off, but *no better bounds are known*. The transition measure is not constant on conjugacy classes, and harmonic analysis seems intractable. A clever coupling may work, but one is not evident. The walk generated by adjacent transpositions (those of the form $(i, i + 1)$), has $|S| = n - 1$, $\Delta = O(n^2)$ (recall “Insertion Sort”), and, thus, a similar $O(n^5 \ln n)$ bound. Again this seems off the mark, but provides the only published bound. \square

Example 2.34 Two-Urn Bernoulli-Laplace. Consider the two-urn Bernoulli-Laplace model with k balls in the left urn and $n - k$ in the right urn. From our earlier examples, this can be viewed as a chain induced symmetrically from an N -bi-invariant transition measure. Without loss assume $k \leq n/2$. We have $\Delta = k$, and $p = 1/2k(n - k)$, and an immediate eigenvalue bound of $\lambda_1 \leq 1 - 1/(32k^3(n - k) + 4k(n - k))$. This is, again, far off the mark. As mentioned earlier, Diaconis and Shahshahani [DS87] get a tight full-spectrum convergence bound and find that the second eigenvalue is exactly $\lambda_1 = 1 - \frac{n}{2k(n - k)}$. We can match their convergence bounds by our earlier couplings for this case. See Theorem 2.10. \square

Example 2.35 Many-Urn Bernoulli-Laplace. For three or more urns, the Bernoulli-Laplace process can still be induced symmetrically from an N -bi-invariant transition measure, but (G, N) is not a Gelfand pair. The harmonic analysis seems intractable. We can still apply Theorem 2.30. For $n = 3k$ balls in $m = 3$ urns, and the strongly-aperiodic form, the diameter of the underlying graph is $2k$ and all non-holding transition probabilities are equal to $p = 1/4k^2$. This gives a bound of $\lambda_1 \leq 1 - \frac{1}{256k^4 + 8k^2}$ or about $1 - \frac{1}{3n^4 + n^2}$. We do better with the coupling bound of Theorem 2.11. \square

Example 2.36 Near-Uniform generation in groups with polynomial diameter. Theorem 2.29 implies that if Γ is a finite group presented as n generators, with a diameter guaranteed bounded by a polynomial $p(n)$, then there is a near-uniform generator for Γ that runs in time polynomial in n and the cost of multiplication in the group. Surprisingly, these ideas can be extended to certain infinite cases. See [Bab90]. \square

Remark: “Reasonable” shuffles. Intuitively, any “reasonable” set of generators of S_n has a Cayley graph with diameter polynomial in n , and the walk on such a graph will therefore have a polynomial convergence guarantee. A good project would be to make these ideas precise. It is known that with probability $3/4$ two randomly chosen elements of S_n will generate the group. If one can get a handle on the distribution of the diameter under two generators, it would be possible to prove

a bound using these theorems. It is hard to conceive of two generators which yield a diameter that is not $O(n^2)$. In a related direction Babai, Kantor, and Lubotsky [BKL89] prove that there exists a constant C such that every non-Abelian finite simple group has a set of generators S , with $|S| \leq 7$, for which the diameter of $\text{Cayley}(\Gamma, \Gamma, S)$ is at most $C \ln |\Gamma|$. \square

Chapter 3

Mean-Value Estimation

Often one wants to determine the mean value

$$h_1 = \sum_{v \in V} h(v)\pi(v) \quad (3.1)$$

of a real-valued function h under some distribution π on a finite set V . If V is very large, or h is sufficiently complex to analyze, computing this value exactly may be prohibitively expensive. In this case one often resorts to estimates obtained by sampling.

If one is able to draw independent samples according to π , one generally makes an estimate using the sample mean

$$A_n = \frac{1}{n} \sum_{i=1}^n h(Y_i), \quad (3.2)$$

where the random variables Y_i are the samples. This is an unbiased estimator, which means that

$$E[A_n] = h_1, \quad (3.3)$$

and its variance is

$$\text{Var}[A_n] = \frac{1}{n} h_2, \quad (3.4)$$

where

$$h_2 = \sum_{v \in V} (h(v) - h_1)^2 \pi(v) \quad (3.5)$$

is the variance of $h(Y)$ when Y is a single element drawn according to π . If bounds on h_2 are known, Chebyshev's inequality may then be used to bound the probability that the estimate lies outside a given interval of h_1 . If $n > \frac{4h_2}{\beta^2}$ then the estimate A_n satisfies

$$\Pr\{|A_n - h_1| \geq \beta\} \leq \frac{1}{4}.$$

The following well-known lemma can then be used to rapidly decrease the probability of error.
(This has been called the 'Powering Lemma' in [JV86].)

Lemma 3.1 (Median Lemma) Let α_i for $1 \leq i \leq m$ be m random estimates of a such that

$$\Pr\{|\alpha_i - a| > \beta\} \leq \frac{1}{4}$$

for each α_i independent of any previous estimates. Let M be the median of the α_i . (If m is even, take the average of the two candidate medians.) Then

$$\Pr\{|M - a| > \beta\} \leq 2^{-m}.$$

Note: The constant $1/4$ may be replaced by any constant $c < 1$, if 2^{-m} is replaced by $c^{m/2}$.

Proof: Since M is the median of the estimates α_i , at least half ($\lceil m/2 \rceil$) of the α_i are at least M and at least half are at most M . Thus if M lies outside the interval $[a - \beta, a + \beta]$ at least $\lceil m/2 \rceil$ of the estimates α_i do. Applying the independence condition with the law of conditional probability, we have $\Pr\{|M - a| > \beta\} \leq (\frac{1}{4})^{\lceil m/2 \rceil} \leq 2^{-m}$. \blacksquare

Thereby $\lceil \frac{4h_1}{\beta^2} \rceil \lceil \lg 1/\delta \rceil$ independent samples are sufficient to get an estimate M of the mean value that is within β of h_1 with probability at least $1 - \delta$. This method of estimating the mean will be called the median-of-sample-means algorithm. Of substantial special interest is the case when h takes values in $\{0, 1\}$ and π is uniform on V . If $H = \{v \mid h(v) = 1\}$, then h is called the indicator function of the set H . The mean value of h is then

$$h_1 = p = \frac{|H|}{|V|}.$$

Typically, this situation comes up in the problem of approximate counting, when one knows $|V|$ and wants to estimate $|H|$, or vice versa. In this case A_n is binomially distributed with parameters n and p . Chernoff-type bounds on the tails of the binomial distribution may be used to get very good bounds on the probability of given error.

These are basic results for mean-value estimation when a source of independent and identically distributed (i.i.d.) samples is readily available. The purpose of this chapter is to provide some analogous analysis of the sample mean and the median of many sample means when the samples are drawn from a time-reversible Markov chain with stationary distribution π on V . The motivation for this investigation is the Markov chain simulation method for sampling. The basic method is to run a Markov chain on V with the desired stationary distribution, running the chain until its distribution is stationary or nearly so, and then drawing samples from the stationary chain. This chapter lays the groundwork for the next chapters, and provides the basic tools for the empirical work in Chapter 5. Here is an outline of the chapter.

Let $\{X_k \mid k \geq 0\}$ be a Markov chain with transition matrix P . To avoid technicalities, we assume X_0 is drawn according to π , so that the chain evolves in the stationary distribution. All of the bounds here hold approximately when X_0 is chosen from a distribution close to π . We will

consider the estimator

$$C_n^t = \frac{1}{n} \sum_{0 \leq i < n} h(X_{it}), \quad (3.6)$$

which is the sample mean when the samples are drawn every t steps in the chain. We will analyze this estimator and estimates based on the median of these sample means.

The following is an easy consequence of the approximation properties of nearly-independent samples discussed in Appendix B.

Theorem 3.2 *Let P be an ergodic Markov chain with stationary distribution π and convergence guarantee $T_P(\epsilon)$. Let X_0 be drawn from any distribution π_0 within ratio $1 + \epsilon$ of π and suppose $t > T_P(\epsilon/2n)$. Then the distribution of the resulting estimator C_n^t is within ratio $(1 + \epsilon)$ of the distribution of the sample mean A_n based on n i.i.d. samples.*

Proof: This follows immediately Theorem B.1 and Theorem B.2. \blacksquare

However, in this chapter we will investigate the performance of the estimator when samples are drawn t steps apart from the stationary (or nearly stationary) chain, where t may be much smaller than the time to stationarity.

For any t , including the case $t = 1$, linearity of expectation still gives

$$\mathbb{E}[C_n^t] = h_1. \quad (3.7)$$

That is, the estimator is unbiased. However, the samples will in general be correlated, and the variance will involve covariance terms.

In Section 3.1 we analyze these covariance terms in terms of the spectrum of the chain. Making a minor extension to a theorem of Aldous, we obtain tight worst-case bounds for the variance of the estimator C_n^t in terms of h_2 , n , t , and λ_1 , the second largest eigenvalue of the chain.

The variance bounds can be used with Chebyshev's inequality to bound the probability that the estimate falls outside a given interval of h_1 . Then the traditional median trick in Lemma 3.1 can be applied to obtain accurate estimates with high probability.

We can also compare the variance of sample means based on truly independent samples, and samples obtained from the stationary Markov chain at various time intervals. The following tradeoff is investigated. Taking t as large as the time to stationarity, gives very 'high quality' samples that are very nearly independent, and the accuracy of C_n^t will nearly match that of independent samples. However, taking t large also means we are expending computational effort to simulate many steps of the chain between samples, but ignoring the 'information' in the intervening states. In particular, when only a λ_1 -based bound on convergence is known, it turns out that it is usually better, from a computational standpoint, to take t small and use this 'information' despite the added correlation effects.

In Section 3.3 we consider the special case when h takes values in $\{0, 1\}$ and π is uniform. For large t , C_n^t is approximately binomially distributed, and Chernoff bounds are available for the sample mean. We discuss the relative merits of close and well-spaced samples when such tight tail bounds are available, and still conclude that using samples separated by only a few steps is typically more efficient than taking nearly independent, well-spaced samples.

3.1 Variance Bounds

In this section we give tight worst-case bounds on the variance of the estimator C_n^t defined in (3.6), when X_0 is drawn from the stationary distribution. In this case the chain continues to evolve in the stationary distribution, and as mentioned in (3.7), C_n^t is an unbiased estimator.

We will be using the notation M , R , B , and Γ and z , introduced in Section 1.2 for the spectral representation of the reversible chain P . The uncertain reader should refer back to that section for definitions. We will not re-define the notation here.

Our analysis closely follows the line of Aldous [Ald87], but we retain some lower order terms in order to obtain a closely matching worst-case lower bound.

We will assume, without loss of generality, that $h_1 = 0$; this assumption does not affect the analysis, but leaves $h_2 = \sum_{v \in V} \pi(v)h^2(v)$ which is simpler to handle.

First write

$$\text{Var}[C_n^t] = \frac{1}{n^2} \left[nh_2 + \sum_{k=1}^n 2(n-k)E_{tk} \right]. \quad (3.8)$$

where

$$E_k = \mathbb{E}[h(X_0)h(X_k)].$$

This is just the standard expansion for the variance of a sum random variables; the first term is the sum of the variances, which may be interpreted as the "independent" component, and the second term is the sum of the covariance terms. To get the form above we have noted that the covariance $\mathbb{E}[h(X_j)h(X_{j+k})]$ between any pair of values $h(X_j)$ and $h(X_{j+k})$ for $0 \leq j \leq n-k$ is the same E_k . A given covariance term E_k appears $2(n-k)$ times, once for each ordered pair at separation k (twice for each unordered pair). We want to bound this sum of covariance terms.

Since X_0 is chosen according to π we get, for any k ,

$$E_k = \sum_{z \in V} \sum_{y \in V} h(z)\pi(z)h(y)P_{zy}^{(k)}. \quad (3.9)$$

Apply Theorem 1.3 to rewrite $P_{zy}^{(k)}$ in its spectral representation, and get

$$E_k = \sum_{z,y} \sum_w h(z)\pi(z)h(y) \sqrt{\frac{\pi(y)}{\pi(z)}} B_{zw}^k \Gamma_{zw} \Gamma_{yw} \quad (3.10)$$

$$= \sum_w B_{ww}^k \left(\sum_{x,y} h(x) \sqrt{\pi(x)} \sqrt{\pi(y)} h(y) \Gamma_{xz} \Gamma_{yw} \right) \quad (3.11)$$

$$= \sum_w B_{ww}^k \left(\sum_x h(x) \sqrt{\pi(x)} \Gamma_{xz} \right)^2 \quad (3.12)$$

$$= \sum_w B_{ww}^k v_w^2. \quad (3.13)$$

In moving between the last two lines we have simply called the the inner sum v_w ; this v_w is just the projection of the vector hR on the basis vector Γ_{zw} in the orthonormal basis of eigenvectors of M .

Since $\Gamma_{zz} = \sqrt{\pi(z)}$ we have

$$v_z = \sum_x h(x) \pi(x) = h_1 = 0 \quad \sum_w v_w^2 = \|hR\|_2^2 = \sum_x \pi(x) h^2(x) = h_2 \quad (3.14)$$

Now one is tempted to bound

$$E_{tk} \leq (\lambda_*)^{tk} h_2. \quad (3.15)$$

This is valid. However, this will yield a bound on the variance of C_n^t in terms of λ_* and thereby lose something. Using an idea of Aldous, we resist this temptation, and obtain a bound in terms of λ_1 only, but only for odd t . The difference is discussed in the remark following the theorem.

Theorem 3.3 (Tight Variance Bounds) *Let $X_0, X_1, \dots, X_{(n-1)t}$ be n states drawn t steps apart from a stationary time-reversible Markov chain P with state space V , stationary distribution π , and second-largest eigenvalue λ_1 . Let h be any real-valued function on V , and define*

$$h_1 = \sum_{v \in V} \pi(v) h(v) \quad \text{and} \quad h_2 = \sum_{v \in V} \pi(v) (h(v) - h_1)^2.$$

If t is odd and positive the mean-value estimator $C_n^t = \frac{1}{n} \sum_{i=0}^{n-1} h(X_{ti})$ has expected value

$$\mathbb{E}[C_n^t] = h_1$$

and variance

$$\text{Var}[C_n^t] \leq \alpha(\tau(t), n) h_2,$$

where

$$\alpha(\tau, n) = \frac{2\tau - 1}{n} + \frac{2\tau(\tau - 1)}{n^2} \left(1 - \left(1 - \frac{1}{\tau} \right)^n \right)$$

and

$$\tau(t) = 1/(1 - \lambda_1^t).$$

Moreover, there exists a function h such that for every t and n ,

$$\text{Var}[C_n^t] \geq \alpha(\tau(t), n) h_2.$$

Proof of the Upper Bound

The expected value is h_1 by linearity. We complete the analysis of the variance from the preceding discussion.

Consider now the sum term from (3.8). In light of (3.13) it is:

$$\begin{aligned} \sum_{k=1}^{n-1} 2(n-k)E_{tk} &= \sum_{k=1}^{n-1} 2(n-k) \sum_w \lambda_w^{tk} v_w^2 \\ &= \sum_w \sum_{k=1}^{n-1} 2(n-k) \lambda_w^{tk} v_w^2 \\ &\leq \sum_w v_w^2 \left(\max_{-1 < \lambda \leq \lambda_1} \sum_{k=1}^{n-1} 2(n-k) \lambda^{tk} \right), \end{aligned}$$

since $v_z = 0$.

Now the maximum of each inner sum is attained at λ_1 . This is because for $\lambda < 0$, and t odd, the inner sum is an alternating sum, and it is bounded above by zero. When λ is positive however, it increases monotonically with λ . So it attains its maximum at $\lambda = \lambda_1$. From this and (3.14), we have

$$\text{Var}[C_n^t] \leq \frac{\Psi}{n^2} h_2 \quad (3.16)$$

where

$$\Psi = n + 2 \sum_{k=1}^{n-1} (n-k) \lambda_1^{tk} = 2 \left(\sum_{k=0}^{n-1} (n-k) \lambda_1^{tk} \right) - n.$$

Putting the sum in closed form, letting $\tau = \tau(t) = 1/(1 - \lambda_1^t)$, and simplifying, gives

$$\begin{aligned} \Psi &= \frac{2}{(1 - \lambda_1^t)^2} \left(\lambda_1^{(n+1)t} - (n+1) \lambda_1^t + n \right) - n \\ &= 2\tau^2 \left[(1 - 1/\tau)^{n+1} + n/\tau - (1 - 1/\tau) \right] - n \\ &= 2\tau^2 \left[(1 - 1/\tau)^{n+1} + n/\tau - (1 - 1/\tau) - n/(2\tau^2) \right] \\ &= 2\tau^2 \left[\frac{n}{\tau} \left(1 - \frac{1}{2\tau} \right) - \left(1 - \frac{1}{\tau} \right) \left(1 - \left(1 - \frac{1}{\tau} \right)^n \right) \right], \end{aligned}$$

so that

$$\frac{\Psi}{n^2} = \frac{2\tau - 1}{n} - \frac{2\tau(\tau - 1)}{n^2} \left(1 - \left(1 - \frac{1}{\tau} \right)^n \right) = \alpha(\tau, n).$$

Then (3.16) gives the desired bound. ■

Proof of the Lower Bound

First, we show how to pick an h such that $E_{tk} = \lambda_1^{tk} h_2$, for every k . Then the result will follow easily.

Let $\Gamma_{z'}$ be an eigenvector of M corresponding to λ_1 . We use $\Gamma_{zz'}$ to denote its z th coordinate. Take the function h to be $h = R^{-1}\Gamma_{z'}$ so that $h(z) = \Gamma_{zz'}/\sqrt{\pi(z)}$. Then, by the orthogonality of the matrix Γ we have

$$\begin{aligned} h_1 &= \sum_z \pi(z)h(z) \\ &= \sum_z \sqrt{\pi(z)}\Gamma_{zz'} \\ &= \Gamma_{z'} \cdot \Gamma_{z'} \\ &= 0 \end{aligned}$$

and

$$\begin{aligned} h_2 &= \sum_z \pi(z)h^2(z) \\ &= \sum_z \pi(z) \frac{(\Gamma_{zz'})^2}{\pi(z)} \\ &= \sum_z (\Gamma_{zz'})^2 \\ &= 1. \end{aligned}$$

By the same orthogonality, this function h satisfies

$$\sum_z h(z)\sqrt{\pi(z)}\Gamma_{zw} = \begin{cases} 1 & \text{if } w = z', \\ 0 & \text{otherwise.} \end{cases}$$

Using this in (3.12) gives

$$E_{tk} = \lambda_1^{tk}.$$

Consequently, by (3.8) this h will attain

$$\text{Var}[C_n^t] = \frac{\Psi}{n^2} h_2 = \alpha(\tau, n)h_2.$$

and the result follows. ■

Remark: If we had used $|E_k| \leq \lambda_1^k h_2$ as in (3.15), we could have avoided the maximization argument in the proof of the upper bound and obtained an upper bound valid for both odd and even t . However, this would have lost something (since $\lambda_1 \leq \lambda_0$) that one cannot recover in a lower bound that is valid for all t . While a lower bound involving λ_0 is possible for even values of t , our lower bound is valid for both odd and even t .

There is a good reason to write the bound in terms of λ_1 rather than λ_0 . Current techniques for bounding the rate to stationarity are based on λ_0 , while most techniques for bounding eigenvalues

can bound only λ_1 . Typically this problem is avoided by converting to a related strongly-aperiodic chain, such as $(I + P)/2$, whose eigenvalues will all be nonnegative. However, this type of conversion may result in a slowdown of a factor of 2. One consequence of this theorem is that when using a chain for estimation we need not make this conversion. \square

Remark: In the case that V is an Abelian group and P a random walk on the group, v_w in (3.13) corresponds to the Fourier coefficient of h at the representation associated with w , and B_{ww} corresponds to the Fourier coefficient of the transition measure at that representation. On groups such as the hypercube (Z_2^d) or the circle (Z_n), where the harmonic analysis is quite tractable and all of the eigenvalues and eigenvectors can be easily written, it may be possible to get better bounds for specific h . For example, if h is such that v_w is known to be small for the large eigenvalues B_{ww} , then a better bound results following essentially the same line of analysis as for Theorem 3.3. However, a natural situation in which such functions h arise is not apparent. \square

We will be interested primarily in the upper bound, which is convenient to apply in the following form.

Corollary 3.4 *With the same setup as in the preceding theorem, consider the estimator C_n^t , taking*

$$n = \lceil (2\tau(t) - 1)k \rceil$$

samples drawn t steps apart, with t odd. Then for every function h and every $k \geq 1$

$$\text{Var}[C_n^t] \leq \frac{h_2}{k}.$$

Proof: Apply Theorem 3.3 with $n = (2\tau - 1)k$. Using the facts that $\tau \geq 1$ and $(1 - \frac{1}{\tau})^n \leq e^{-n/\tau} \leq \frac{1}{2}$ for the given n , we get

$$\begin{aligned} \alpha(\tau, n) &= \frac{2\tau - 1}{n} - \frac{2\tau(\tau - 1)}{n^2} \left(1 - (1 - \frac{1}{\tau})^n \right) \\ &\leq \frac{2\tau - 1}{n} - \frac{2\tau(\tau - 1)}{n^2} (1 - e^{-n/\tau}) \\ &\leq \frac{1}{k} - \frac{\tau(\tau - 1)}{(2\tau - 1)^2 k^2} \\ &\leq \frac{1}{k}. \end{aligned}$$

Note: for the typical case when $\tau > 1$, the inequality is, in fact, strict. \blacksquare

The interpretation of the corollary above is that taking $\lceil (2\tau(t) - 1)k \rceil$ samples drawn t steps apart from the stationary Markov chain gives an estimator whose variance is guaranteed to be at most that of the sample mean based on k independent samples.

Applying Chebyshev's inequality with the preceding corollary, immediately gives the following confidence interval bound.

Corollary 3.5 (Chebysheff Bounds) For $n \geq (2\tau(t) - 1)k$, and t odd, we can guarantee

$$\Pr\{|C_n^t - h_1| > \beta\} \leq \frac{h_2}{k\beta^2}.$$

In particular if $n \geq (8\tau(t) - 4)h_2/\beta^2$, then

$$\Pr\{|C_n^t - h_1| > \beta\} \leq \frac{1}{4}.$$

Repeated estimates C_n^t can be combined using the median trick of Lemma 3.1, provided that each estimate satisfies the bound of Corollary 3.5 independent of the result of preceding estimates. This can be guaranteed by drawing the initial sampling point X_0 for each estimate independently, or even nearly independently. This is discussed further in Appendix B. (In Chapter 4 we even show a case in which one can use repeated estimates when the points X_0 are highly correlated, but still have certain "sample majority" properties resembling those of independent samples.)

Remark: One interpretation of Corollary 3.5 is that the estimator C_n^t based on $[(2\tau(t) - 1)k]$ samples is essentially as good as the sample mean based on k pairwise-independent π -distributed samples.

To apply our bounds in practice we rely on an upper bound on h_2 . If h is a member of a class of functions where h_2 can be bounded *a priori*, our bounds may be applied directly. For example, when h is the indicator function of a subset $H \subseteq V$, and π is the uniform distribution, we have $h_2 \leq 1/4$. In general, however, one may have to resort to estimates of h_2 as well. Various estimators for h_2 for general stationary regenerative processes have been suggested [Han57] [Igl78] [GI84]. Glynn and Iglehart [GI87] give a joint central limit theorem for a mean and variance estimator, and Calvin [Cal90] provides some additional analysis but not any non-asymptotic bounds. \square

3.2 Sampling to Achieve Given Variance

The point of this section is to show how the bounds of the previous sections can be applied to answer a basic question that arises when computing mean values by the Markov chain simulation method. Namely, we discuss the question 'How large should one choose the spacing t between samples in order to obtain estimates with a given guaranteed variance bound?' The answer to this question will also answer similar questions which arise when applying the median-of-sample-means method with Chebysheff bounds.

Suppose that we want to estimate h_1 with variance about h_2/k (the variance of the sample mean based on k independent samples). For the moment, let us restrict our discussion to the typical case. Suppose we have found a bound on λ_1 . After conversion to the aperiodic chain $P' = I + P$, one gets a convergence guarantee for P' whose leading term is $\frac{1}{1-\lambda_1} \ln(1/\pi_{\min})$ or about $T = 2\tau(1) \ln(1/\pi_{\min}) \geq 2\tau(1) \ln |V|$, where $\tau(t)$ is that of the original chain P .

By taking a small t , we use much of the 'information' in the states that we pass through. However, we need more samples to achieve a given variance, and each time we take a sample we must compute the function h . By taking a large t , we get nearer to independent samples and reduce the number of samples required. We compute the function h fewer times. In tradeoff we spend time to simulate many steps of the chain between samples.

For example, suppose we take $t > T$. These samples are roughly independent, and perform about as well in variance. (This is by Theorems B.1 and B.2, or by Corollary 3.4 since $\tau(t) \approx 1$ for $n > T$.) This will mean that we compute h only about k times to get variance h_2/k , but that we run the chain on the order of kT steps.

Alternatively, we might take $t = 1$, so that we are taking successive samples. Then $n = (2\tau(1) - 1)k$ samples will be enough to guarantee a variance at most h_2/k . So h will be computed $n = (2\tau(1) - 1)k$ times. But on the other hand, we will only need to run the chain for $T + (k - 1)(2\tau(1) - 1)$ steps (T to get the initial sample X_0 and $(k - 1)(2\tau(1) - 1)$ to get the remaining samples). This is much smaller than kT if the state space V is large. Thus, if the cost of computing h is negligible compared to the cost of taking a step of the Markov chain, *taking successive samples will be cheaper than taking roughly-independent samples to achieve a given variance*.

We should also consider the possibility that h may be calculated more easily incrementally than from scratch. In other words, it may be that computing $h(X_i)$ from scratch requires significant resources, while it may be easy to compute $h(X_i)$ when we are given $h(X_{i-1})$, and some information about the transition that took X_{i-1} to X_i . For example, suppose $v = (v_1, v_2, \dots, v_n)$ is an integer vector in $[m]^n$, and that $h(v) = \sum_{j=1}^n v_j^2$. This requires time linear in n to compute in general. Suppose, however, that we sample using a chain generated by moves of the following type: if $X_i = (v_1, v_2, \dots, v_n)$ is the current position, then uniformly choose a random step direction $\sigma \in \{+1, -1\}$, and a random coordinate $j \in [n]$; then, if $(v_j + \sigma) \in [m]$, move by adding σ to v_j ; otherwise, remain at the current position. (This move generates a symmetric ergodic chain. Spectral analysis of this type of chain is in Chapter 2.) For this chain we can compute $h(X_i)$ in constant time from $h(X_{i-1})$ and the knowledge of the move (σ, j) that took X_{i-1} to X_i , because $h(X_i) = h(X_{i-1}) + 2\sigma v_j + 1$. A similar trick applies to the computation of the χ^2 statistic on tables that we investigate in Chapter 5.

Of course a spacing t somewhere between 1 and T may be best. To illustrate, suppose we have a function h which is fairly expensive to compute, requiring about the same cost as about $\tau(1)$ steps of the chain. Take an odd t near $2\tau(1)$. This makes $(2\tau(t) - 1) \approx 1.3$.¹ Now if we let $n \approx 1.3k$, computing C_n^t will take $T + 2.6\tau(1)(k - 1)$ steps of the chain, but has the additional property that we need only compute the function h about $1.3k$ times. The total computational cost will be the equivalent of about $T + 4k\tau$ steps of the chain. If τ is large, this is considerably less than the approximately $T + 2\tau^2k$ steps needed when successive sampling, and is also less than the $(T + \tau)k$ steps needed for nearly independent samples drawn at spacing T , given the reasonable assumption

¹ Assuming that λ_1 is very near 1, we have $2\tau(2\tau(1)) - 1 \approx (2/(1 - \epsilon^2)) - 1 \approx 1.3$.

that $T > 3\tau$.

In general it seems that to achieve a mean value estimate of given variance, when only a bound on λ_1 is known, it is typically better to take samples spaced at a distance significantly smaller than the time required to reach stationarity. This is somewhat surprising. The best results in this direction come when the walk is on a family of (d, c) -magnifiers. (For a definition, see Section 1.5.2.) We explore the consequences in the next chapter.

Here we illustrate with some simpler examples.

Example 3.6 (Random Walk on the Hypercube) Let P be the natural aperiodic random walk on the additive group (\mathbb{Z}_2^d, \oplus) , (the hypercube). This walk is generated by repeating the following step: A coordinate i is chosen uniformly from $\{0, 1, 2, \dots, d\}$. If $v \in \mathbb{Z}_2^d$ denotes the current position, the new position is $v \oplus e_i$, where $e_0 = (0, 0, 0, \dots, 0)$, and for $j \neq 0$, e_j has a single 1 in the j th coordinate. The chain is symmetric, and the stationary distribution is uniform on \mathbb{Z}_2^d . The entire spectrum of the chain is easily obtained by harmonic analysis [Dia88], or using the construction of the hypercube as the Cartesian product of 2-cliques (Chapter 2). The eigenvalues are the values $1 - \frac{2j}{d+1}$ ($j = 0, 1, \dots, d$) appearing with the respective multiplicities $\binom{d}{j}$. A λ_1 bound alone tells us $T = O(d^2)$ steps are sufficient to get roughly independent samples. Using the full spectrum one can show that $\frac{1}{4}d(\ln d + c)$ steps are sufficient to get samples within total squared variation $\frac{1}{2}(\exp(e^{-c}) - 1)$, and the leading term $\frac{1}{4}d \ln d$ is essentially tight. (See [Dia88].) However, $2\tau(1) = (d+1)$, so that taking kd adjacent samples instead of k samples spaced $d \ln d$ steps apart gives an estimate with as small a variance, and saves steps of the chain. This is an artificial example, because in practice we would not use this walk for sampling from the set of binary d -tuples. However it illustrates that the theorem may yield a savings for sampling from the chain, even when a tight full-spectrum bound is known.

□

Example 3.7 (Permutations via Random Transpositions) The following is a well-known linear time algorithm for constructing a uniform random permutation of the set $[n]$. (It is attributed by Knuth [Knu68, Vol. 2, Sec. 3.4.2, p. 140] to L. E. Moses and R. V. Oxford [MO63].) Assume that $\text{random}(i, n)$ returns a random integer in the set $\{i, i+1, \dots, n\}$ in constant time.

Random Permutation

```

for 1 ≤ i ≤ n do p[i] ← i
for 1 ≤ i ≤ n - 1 do begin
    j ← random(i, n)
    swap p[i] and p[j]
end

```

This yields a random permutation in the array p in time $\Theta(n)$. The algorithm uses $\Theta(n \ln n)$ random bits.

Alternatively, we can generate a random permutation by repeated random transpositions, using the Markov chain generated by taking moves of the following type: uniformly choose an unordered pair $\{i, j\} \in [n]^2$, (also allowing $i = j$); then swap $p[i]$ and $p[j]$. From the analysis of this walk in [DS81] and [Dia88] we know that essentially $\frac{1}{2}n \ln n$ such operations are both sufficient and necessary to make the permutation nearly uniformly distributed. So it would seem that the random walk method provides a strictly inferior method of sampling permutations. But this is not necessarily the case when sampling for mean-value estimation.

From [DS81] we know that $\lambda_1 = 1 - \frac{2}{n}$ (exactly) for this chain, and thus that $2\tau(1) = n$. Thus starting at an initial random permutation (which we can generate by the traditional method), in kn steps of the walk we can get a mean-value estimate (based on kn samples) with variance *slightly smaller* than the variance of k real uniform samples. This method of using the random walk takes $\Theta(kn)$ time and $\Theta(kn \ln n)$ random bits, for both resources the same order as required to build the k independent samples. So the random walk method here is competitive. The winning method will be determined by implementation-dependent constants. In fact, if h is easily computed incrementally, the random walk method may be more efficient than this traditional algorithm for the problem. \square

Example 3.8 (Random Subsets via Bernoulli-Laplace) A similar situation occurs when generating random subsets. R. W. Floyd [Flo] has invented the following beautiful algorithm for producing a subset S of cardinality m from the set $[n]$, distributed uniformly at random over the $\binom{n}{m}$ possibilities. Again, assume $\text{random}(i, n)$ returns a random integer in the set $\{i, i+1, \dots, n\}$ in constant time.

Random Subset

```

 $S \leftarrow \emptyset$ 
for  $j \leftarrow k$  downto 1 do begin
     $r \leftarrow \text{random}(j, n);$ 
    if  $r \notin S$  then insert  $r$  in  $S$  else insert  $j$  in  $S$ 
end

```

Assume that n is even and $m = n/2$ (any value $\Theta(n)$ will do). Then, if we represent the set S using a bit string (a string of n bits with bit s equal to 1 if $s \in S$, and 0 otherwise), the algorithm runs in time $\Theta(n)$ and requires $\Theta(n \lg n)$ random bits.

The two-urn Bernoulli-Laplace process, for which we gave a coupling analysis in Chapter 2, yields a nearly-random permutation in $\Theta(n \ln n)$ steps. Again, while it seems at first that the random walk method is strictly inferior, this is not necessarily the case when sampling for mean-value estimation.

For this process we know $\lambda_1 = 1 - \frac{n}{m(n-m)}$ [DS87]. Thus, for $m = n/2$, we have $2\tau(1) = n/2$. Starting from a uniform random subset (which we can generate by Floyd's algorithm), and successively sampling over an $kn/2$ steps of the walk we can get an unbiased mean value estimate with variance

at most h_2/k . This means the random walk method may be competitive with Floyd's algorithm. Both methods take time $\Theta(kn)$ and use $\Theta(kn \ln n)$ random bits. Again, the random walk method may perform better if h is easily computed incrementally. \square

Example 3.9 (Estimation on Random Matchings) Many interesting statistical and combinatorial sample spaces may be associated with sets of permutations where the image of each position is restricted to lie within a given set, dependent upon that position. Such spaces are naturally identified with the set of matchings (1-factors) of a bipartite graph. The counting problem for such sets is provably hard ($\#P$ -complete), so there are good reasons to believe that traditional methods of exact uniform sampling in this space will not be fruitful. (See [JV86].) Broder [Bro86] suggested a Markov chain for near-uniform sampling from this set. Jerrum and Sinclair [JS88] gave a system of canonical paths for this chain. They showed that when the underlying bipartite graph is sufficiently dense, their Cheeger-type bound using these paths yielded a bound of $\lambda_1 = 1 - \Omega(n^{-12})$, implying nearly-independent samples could be obtained at a cost of $O(n^{13} \ln n)$ steps per sample. Recently J. A. Fill used the same paths in a Poincaré-type bound yielding $\lambda_1 \leq 1 - \Omega(n^{-7})$. (This calculation appears in [DS89].) This implies that nearly independent samples can be obtained at a cost of $O(n^8 \ln n)$ steps of the chain. For estimation of mean-values on the space, however, this means that $O(n^7)$ steps per sample will be enough to yield estimates with equally small variance. This combination of results makes a big difference in the range of n for which any practical means of accurate estimation is currently known. A substantial improvement also applies to mean-value estimation using Jerrum and Sinclair's "all-matchings" chain. \square

3.3 Indicator Functions

We now consider the special case where h is the indicator function of a set H and π is the uniform distribution. Recall that h is called the indicator function of the subset H if $h : V \rightarrow \{0, 1\}$ and $H = \{v \mid h(v) = 1\}$. (Mathematicians would usually refer to these as characteristic functions, a term which we avoid because it has a different meaning for statisticians.) The mean value of h under π is then $h_1 = p = \frac{|H|}{|V|}$. Estimating the mean value of an indicator function turns out to be the case of principal interest in most practical problems; one wants to estimate the proportion of elements of a set that satisfy a given criterion. Indicator functions also arise when one estimates a probability density of a random variable by sampling. One partitions the range of the variable into a number of subranges called 'bins.' Then one counts how many samples fall into a given bin. The graph of the result is called a histogram. Estimating the density in this way may be viewed as simultaneously estimating the mean value of the indicator function associated with each bin.

If we draw samples Y_1, Y_2, \dots, Y_n independently and uniformly from V then each value $h(Y_i)$ is a Bernoulli random variable with parameter $h_1 = p$. The single-sample variance is $h_2 = p(1-p) \leq 1/4$. The sum $\sum_{i=1}^n h(Y_i)$ is distributed Binomial(n, p), and the following well-known Chernoff-type bound

provides a means of bounding the relative error in the estimator A_n of (3.2).

Lemma 3.10 *If $X \sim \text{Binomial}(n, p)$, and $q = 1 - p$ then*

$$\Pr\left\{\frac{X}{n} - p > \beta\right\} \leq e^{-\beta^2 n / 4pq} \quad (3.17)$$

$$\Pr\left\{\frac{X}{n} + p < \beta\right\} \leq e^{-\beta^2 n / 4pq} \quad (3.18)$$

Proof: See, for example, [CLR90, Chapter 6]. \blacksquare

From Theorem 3.2, we know that the distribution of C_n^t for large t must also be approximately $\text{Binomial}(n, p)$. Combining this with Lemma 3.10 above, we get the following theorem.

Theorem 3.11 (Chernoff Bound) *Let $\{X_k \mid k \geq 0\}$ be a time-reversible Markov chain with transition matrix P , stationary distribution π , and convergence guarantee $T_P(\epsilon)$. Let $t > T_P(\epsilon/2n)$. If h is the indicator function of any subset $H \subseteq V$, and*

$$C_n^t = \frac{1}{n} \sum_{i=1}^n h(X_{it})$$

then

$$\Pr\{|C_n^t - p| > \beta\} \leq 2(1 + \epsilon)e^{-\beta^2 n / 4pq},$$

where $p = h_1 = \frac{|H|}{|V|}$ and $q = 1 - p$.

Proof: The result follows by applying the preceding lemma, Theorem 3.2 and the sum rule to each tail. \blacksquare

The bounds on the sample means obtained in this way are certainly stronger than Chebysheff bounds. However, for mean-value estimation, they do not represent a significant improvement over the combination of the Chebysheff bounds and the median-of-sample-means method.

Consider these two alternatives: (1) Take well-spaced samples and compute the sample mean, bounding the error using the preceding Chernoff bounds. (2) Take closely spaced samples and use the Chebysheff bounds of Corollary 3.5 together with the median-of-sample-means method to reduce the probability of error. In either case, $\Theta((pq/\beta^2) \ln(1/\delta))$ samples will be necessary and sufficient to obtain an estimate within β of the true mean with probability at least $1 - \delta$. However, using the Chebysheff and median method we will typically be able to achieve the desired accuracy with substantially fewer steps of the chain.

Let $T = 2\tau(1) \ln |V|$, as in Section 3.2. Once again, T is essentially the leading term in our usual second-eigenvalue bound on the time to stationarity. From Corollary 3.5 and using the fact that $h_2 \leq 1/4$, we know that $(2\tau(1) - 1)/\beta^2$ successive samples from the stationary chain are sufficient to guarantee that $\Pr\{|C_n^t - p| > \beta\} \leq 1/4$ for any indicator function h . Thus essentially

$$(T + (2\tau(1) - 1)/\beta^2) \lg(1/\delta)$$

steps will allow us to get a median-of-sample-means estimate M that satisfies

$$\Pr\{|C_n^t - p| > \beta\} \leq \delta.$$

Using well-spaced samples ($t \geq T$) we would need about $T/\beta^2 \lg(1/\delta)$ steps. So, again, taking successive samples saves us a factor of about $\ln|V|$ in the number of steps of the chain that we need to simulate.

Remark: The other issues discussed in Section 3.2 concerning the cost of computing h and the possibility of computing h incrementally are pertinent here as well and should be considered when making a choice of t . \square

3.4 Central Limit Theorem

In the case of independent samples, the Central Limit Theorem tells us that in the limit as n approaches infinity, the variable

$$Z_n = \frac{1}{\sqrt{n}}(A_n - h_1)$$

has the normal distribution with mean 0 and variance 1, and tighter non-asymptotic bounds are possible with appeal to theorems of the Berry-Esséen type [Hal82].

There are various versions of the Central Limit Theorem that also hold for Markov chains. These say that the distribution of the sample mean of any function on the state space converges to a normal distribution. In the present environment we have the following asymptotic version.

Theorem 3.12 For any $t \geq 1$, the estimator C_n^t with X_0 drawn according to π , is asymptotically normally distributed with mean h_1 and variance $\text{Var}[C_n^t]$. More precisely, for each fixed z

$$\lim_{n \rightarrow \infty} \Pr\left\{\frac{C_n^t - h_1}{\sqrt{\text{Var}[C_n^t]}} \leq z\right\} = \Phi(z),$$

where

$$\Phi(z) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^z e^{-u^2/2} du$$

is the cumulative distribution of a standard Normal random variable.

Proof: This is the combination of Theorems 1, p. 99, and Theorem 3, from [Chu60]. \square

This suggests approximating the distribution of C_n^t with the Normal distribution with mean h_1 and variance $\text{Var}[C_n^t]$. We can combine this idea with our bounds on the variance of C_n^t , and use the fact that a given confidence interval around the mean of a Normal distribution shrinks monotonically with shrinking variance. This motivates the following approximate confidence intervals. We emphasize that these are only approximate; they provide some intuition, but not much assurance.

Consider the estimator C_n^t with t odd, and X_0 drawn according to π . For n large and at least $2\tau(t) - 1$ one has

$$\Pr\{|C_n^t - h_1| \geq \beta\sqrt{h_2/n}\} \text{ is approximately bounded by } 2\Phi(-\beta).$$

To illustrate, taking any $t \geq 1$ and n large, (at least large enough that $n \geq \max\{2\tau(t) - 1, 660h_2\}$), should give $C_n^t = h_1 \pm .01$ with approximately 99% confidence. (The number $660h_2$ is from [Fel70, Vol. 1, p. 245].) That is,

$$\Pr\{|C_n^t - h_1| \geq .01\} \text{ is approximately bounded by } .01.$$

Remark: An interesting avenue for further investigation would be to research what is known about non-asymptotic versions of the Central Limit Theorem for Markov chains. Likely there will be a spectral formulation, and the new eigenvalue bounds from Chapter 1 will find additional applications there.

Certainly, for large sample spacings t one could apply Theorem 3.2 with known eigenvalue bounds and known non-asymptotic versions of the central limit theorem for n i.i.d. variables, (see [Hal82]), to get rigorous non-asymptotic results on the distribution of C_n^t . We will not investigate these topics here.

Using the same techniques as in Theorem 3.3, bounds for the higher moments of sample means from the stationary chain are also feasible. However a notable consequence of the discussion in Section 3.3 is that for the large class of practical problems that require mean-value estimates for indicator functions, it will not yield substantial payoffs to seek higher moment bounds. \square

Chapter 4

Using Expanders in Estimation

Interesting consequences arise when the variance bounds of the preceding chapter are applied to random walks on expander graphs.

A family of graphs \mathcal{G} is a family of degree- d ϵ -enlargers if for each $G \in \mathcal{G}$, G has degree bounded by d and the first positive eigenvalue ν_1 of the graphical Laplacian $Q(G) = D - A$ satisfies $\nu_1 \geq \epsilon$, for a fixed $\epsilon > 0$ independent of G . Alon [Alo86] showed that this is equivalent to guaranteeing that \mathcal{G} is a family of magnifiers. (In Section 1.5.2, we discussed this result in the Markov chain setting.) Such expanding families have been studied by several authors, and they have numerous applications. A number of explicit constructions of infinite families are now known. (See e.g. [GG81], [JM85], [LPS86], [AGM87].)

Theorem 4.1 (Expander-based Estimates) *Let \mathcal{G} be a family of d -regular ϵ -enlargers that are not bipartite. Let $G = (V, E)$ be any graph in \mathcal{G} . Let C_n^t be the estimator defined in (3.6) where P is the natural random walk on G and X_0 is drawn according to the uniform stationary distribution, and define*

$$\tau(t) = \frac{1}{1 - (1 - \epsilon/d)^t}.$$

For any function h on V , and any odd t , and any $k \geq 1$ if $n \geq (2\tau(t) - 1)k$, then we have

$$\mathbb{E}[C_n^t] = h_1,$$

and

$$\text{Var}[C_n^t] \leq \frac{1}{k} h_2.$$

Proof: The natural random walk P on any member of a family of d -regular ϵ -enlargers has $\lambda_1(P) \leq 1 - \epsilon/d$. This follows immediately from the fact that $L(P) = \frac{1}{d}Q(G)$, and $\lambda_1 = 1 - \mu_1 = 1 - \nu_1/d$. (See Section 1.5.1.) Thus for this chain the function $\tau(t) = 1/(1 - \lambda_1^t)$ is identical to the function $\tau(t)$ stated in the theorem. Combining this with Corollary 3.4 yields the result immediately. \blacksquare

Note that for all t , we have $\tau(t) \leq \tau(1) = d/\epsilon$, which is a constant independent of the size of the vertex set. So roughly speaking, this theorem says that given one element drawn uniformly at random from V , we can get pretty good additional samples (for mean-value estimation) at the cost of a constant number of steps per sample, where this constant does not depend on $|V|$. Moreover, the fact that the family is degree-bounded suggests that we may be able to take each step of our walk using a constant number of random input bits. In this case we will be able to good estimates at a random-bit cost *much smaller* than the cost of independent samples.

In this chapter we will show a result of this type, and how to merge it with the techniques in [AKS87, IZ89, CW89] to produce an algorithm for estimating the mean value of a function $h : \{0, 1\}^n \rightarrow \mathbb{R}$ using very few random bits.

If h_1 is the mean and h_2 is the variance of h under the uniform distribution, the algorithm uses $n + O(h_2/\beta^2) + O(\lg(1/\delta))$ independent random bits to produce $O((h_2/\beta^2)\lg(1/\delta))$ samples in $\{0, 1\}^n$. These samples are used in the standard median-of-sample means algorithm, to obtain an estimate M such that

$$\Pr\{|M - h_1| > \beta\} \leq \delta.$$

The scheme presented here is naturally viewed as a pseudo-random generation scheme for the median-of-sample-means algorithm, which usually requires $\Theta(n(h_2/\beta^2)\lg(1/\delta))$ independent random bits to achieve the same error bounds.

The method suggested here is ultimately not optimal at saving random bits, but this chapter is intended to serve as an example of the application of the results of the last chapter, and the use of expanders for sampling. The main new result is Theorem 4.3.

Our algorithm uses random walks on expander graphs in a combination of two separate stages of pseudo-random generation. The techniques of [AKS87] in [IZ89] and [CW89] do not alone yield savings as good as we present here. We discuss the relationship to these and other similar results in Section 4.7.

4.1 Preliminaries

We will assume the reader is already familiar with the notation introduced in Chapters 1 and 3.

Let $V = \{0, 1\}^n$, where n is even. (In practice, if n is odd, one can 'pad' the domain V by one bit to make n even.) Consider a function $h : V \rightarrow \mathbb{R}$. The mean value h_1 under the uniform distribution V is

$$h_1 = \frac{1}{2^n} \sum_{v \in V} h(v).$$

Define

$$h_2 = \frac{1}{2^n} \sum_{v \in V} (h(v) - h_1)^2.$$

This is the variance of $h(X)$ when X is an element chosen uniformly at random from V . We will assume throughout that h_2 is known, or at least that an upper bound is known. In the remainder of this chapter h_2 may generally be replaced where it appears by such an upper bound. Note that whenever h is an indicator function (i.e., $h : V \rightarrow \{0, 1\}$), we have $h_2 \leq \frac{1}{4}$, since it is equal to the variance of a Bernoulli random variable.

We will be applying the following magnifier introduced by Alon, Galil, and Milman [AGM87]. The graph has vertex set $Z_m \times Z_m$, for arbitrary nonnegative integer m . On this vertex set define the permutations

$$\begin{aligned}\sigma_0(x, y) &= (x, y) \\ \sigma_1(x, y) &= (x, y + 2x) \\ \sigma_2(x, y) &= (x, y + 2x + 1) \\ \sigma_3(x, y) &= (x + 2y, y) \\ \sigma_4(x, y) &= (x + 2y + 1, y)\end{aligned}$$

where all additions are modulo m . The graph is obtained by connecting every vertex $v = (x, y)$ to the 9 vertices consisting of $\sigma_i(x, y)$ and the inverse images $\sigma_i^{-1}(x, y)$, $0 \leq i \leq 4$. (Note two things: (1) $\sigma_0 = \sigma_0^{-1}$ is the identity map, and (2) we use the term graph in our general sense, as defined on p. xiii.) For each (even) n , let G_n be the graph obtained by letting $m = 2^{n/2}$ in the above. Associate with each vertex (x, y) the binary n -tuple corresponding to the concatenation of the binary representations of x and y . In this way, we may view G_n as a magnifier with 2^n vertices labelled with the elements of $V = \{0, 1\}^n$.

For every n , the graph G_n has Laplacian eigenvalue $\nu_1 \geq 8 - 5\sqrt{2}$, independent of n . (See [JM85] and [AGM87, p. 341].) This implies that if P is the *strongly aperiodic* form of the random walk process on G_n , then $\lambda_n(P) = \lambda_1(P) \leq 1 - (8 - 5\sqrt{2})/18 \leq .9485$ for every n . (We work here with the *strongly aperiodic* form of the random walk on G . For background see Section 1.4.2.) Let

$$t = 23.$$

This value is large enough so that

$$\lambda_1^t \leq 1/10,$$

and for this t we have

$$\tau(t) \leq 10/9,$$

where $\tau(t) = \frac{1}{1-\lambda_1^t}$ is the same function as that used in Chapter 3.

In addition to t , the following parameters are used later. For now, these should simply be remembered to be constants. The interpretation of these constants will be given when they appear,

but their precise values are not crucial to the arguments.

$$\begin{aligned}c_1 &= 2 \\c_2 &= 225 = \lceil 8t\tau(t) \rceil \\c_3 &= 736 = 8t\lceil \lg 9 \rceil.\end{aligned}$$

Remark: The exact choice of magnifier is not crucial for our purposes. However, the following properties are used and the reader should notice that the graphs G_n defined above have these properties. The graphs should be connected in order to guarantee that the random walk on the graph is irreducible. The graphs should be regular, with degree bounded by a constant in n . Also, the graphs should have a compact description so that if we know our current position, the choice of one of the edges and its traversal in the walk can be simulated by an efficient computation on the binary n -tuple corresponding to the current vertex. Also desirable is that the eigenvalue λ_1 be bounded well away from 1, since this increases the efficiency of the method. \square

4.2 Outline of the Algorithm

Our algorithm for mean-value estimation consists of three stages which we first summarise. The first two stages together produce a set of samples drawn from V . The final stage produces the estimate by using the samples in the basic median-of-sample-means algorithm described in Chapter 3. A precise description of the new algorithm is given in Section 4.6. However, it will not be very readable without the background provided by the intervening sections. The illustration in Figure 4.2 on the next page may be helpful in interpreting the initial summary here.

In the first stage we use $s + O(\lg(1/\delta))$ random input bits in a random walk on G_s , where $s = n + O(h_2/\beta^2)$. This walk is used to generate $O(\lg(1/\delta))$ (correlated) random strings of length s . These we call sampling seeds.

In the second stage, each of these $O(\lg(1/\delta))$ sampling seeds is used as the random input needed to specify another walk, but this time on the graph G_n . This random walk consists of $O(h_2/\beta^2)$ steps. Recall that the vertex set of G_n is $V = \{0, 1\}^n$, so this is the sample space from which we want to draw our samples, and we do so during this walk.

Finally, the samples obtained from each sampling seed are used to produce a sample mean, by evaluating h at each sample and averaging. The final estimate M is the median of the resulting $O(\lg(1/\delta))$ sample means. This M will lie within β of h_1 with probability at least $1 - \delta$.

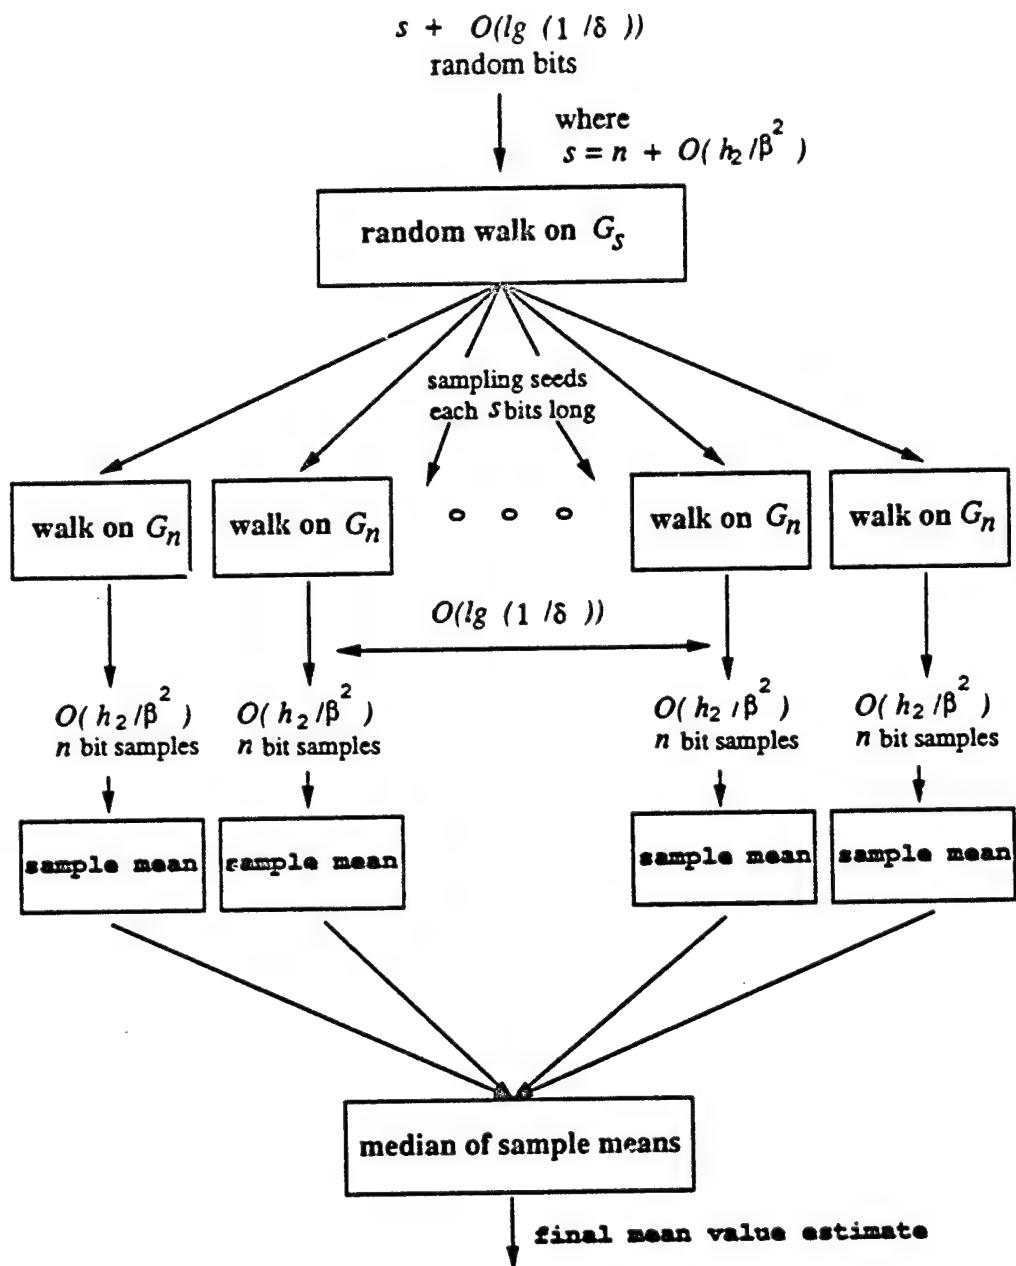


Figure 4.2: Schematic Outline of the Estimation Algorithm.

4.3 Sample Means from G_n

The following theorem is one of the two results forming the basis for our algorithm. It shows, roughly speaking, that in computing the sample mean of any real-valued function h on $V = \{0, 1\}^n$, we can replace k independent samples from V by $k' = O(k)$ correlated samples which we obtain with a number of random bits much smaller than nk (the entropy of the k independent samples from V). Moreover, this substitution does not increase the variance of the sample mean.

Theorem 4.3 (Expander Sample Means) *Let c_1, c_2 , and t be as defined in Section 4.1. There is a deterministic algorithm that given only $n + c_2 k$ independent uniform random bits, outputs $k' \leq c_1 k$ random binary n -tuples $X_1, X_2, \dots, X_{k'}$ with the following properties.*

If h is any real-valued function on $\{0, 1\}^n$, let $S_X = \frac{1}{k'} \sum_{i=1}^{k'} h(X_i)$, and let $S_Y = \frac{1}{k} \sum_{i=1}^k h(Y_i)$, where Y_1, Y_2, \dots, Y_k are independent uniform binary n -tuples.

Then for every h we have,

$$\mathbb{E}[S_X] = \mathbb{E}[S_Y] = h_1$$

and

$$\text{Var}[S_X] \leq \text{Var}[S_Y] = h_2/k.$$

Proof: We generate the samples X_i using the strongly aperiodic random walk process P on the graph G_n above, using the random bits to make the necessary choices. Use the initial n bits as a binary n -tuple specifying the initial vertex of the walk. This starts us in the uniform stationary distribution. We draw the samples X_i every t steps apart from this stationary Markov chain. Now Theorem 3.3 tells us that the sample mean based on $k' = \lceil (2\tau(t) - 1)k \rceil \leq c_1 k$ such samples has the stated properties. We can perform the strongly aperiodic walk on G_n using only $4 = \lceil \lg 9 \rceil$ random bits per step, so $n + 4t\lceil (2\tau(t) - 1)k \rceil \leq n + c_2 k$ random bits are sufficient to take this walk. This gives the result. ■■

Remark: The constants c_1 and c_2 depend on the choice of t set in Section 4.1. Larger t will bring c_2 close to 1, but make c_1 larger.

Note that only a constant number of random bits are required to take each step of the walk on G_n , but each step requires $\Theta(n)$ time in general. If random bits are available at unit cost, the time complexity of using nk random bits or taking $k' = O(k)$ steps of the walk are the same $\Theta(nk)$. Here we are concerned with saving random bits, and incurring only small additional time costs. □

Corollary 4.4 *There is a deterministic algorithm that given $\beta > 0$ and $n + c_2 \lceil 100h_2/\beta^2 \rceil$ uniform independent random bits, outputs an estimate S satisfying*

$$\Pr\{|S - h_1| > \beta\} \leq \frac{1}{100}.$$

The algorithm runs in time that is linear in n , h_2 , and the cost of computing h , and quadratic in $1/\beta$.

Proof: The estimate S is simply the sample mean S_X based on the samples obtained by the preceding theorem. The inequality is Chebyshev's. The running time bound is clear. \blacksquare

To get an estimate that is within an interval of width β around h_1 with probability at least $1 - \delta$, we could now simply apply Lemma 3.1 and take the median of $O(\lg(1/\delta))$ independent estimates S of this type. However, this would use $\Theta((n + h_2/\beta^2)\lg(1/\delta))$ random bits, and we don't want to incur that cost. The results of the next section suggest how we can get away with using only $n + O(h_2/\beta^2) + O(\lg(1/\delta))$ random bits.

4.4 Majorities from G_n

We have already seen that $(2\tau(t) - 1)k$ samples drawn a constant t steps apart from G_n have approximation properties similar to k i.i.d. samples in the sense that sample means based on $O(k)$ samples drawn from G_n have variance matching k i.i.d. samples. We now show how the same type of samples also have properties similar to iid samples under certain majority tests.

The results of this section are really those of [AKS87] [IZ89] and [CW89] translated slightly. We will show that if we consider any small subset B of the vertices, not too many of the samples we draw will lie in B . This provides the second key ingredient for the algorithm.

Define the projection matrix N of a subset $B \subset V$ as the matrix $|V| \times |V|$ indexed by the elements of V such that $N_{vv} = 1$ if $v \in B$ and all other entries of N are zero.

For vectors $\phi \in \mathbb{R}^{|V|}$, let $|\phi|_2 = (\sum_{v \in V} \phi^2(v))^{1/2}$ denote the L^2 norm, and let $|\phi|_1 = \sum_v |\phi(v)|$ denote the L^1 norm.

We begin with the following lemma.

Lemma 4.5 (From [IZ89]) Let P be the strongly aperiodic walk on G_n , and let t be as defined in Section 4.1. Let $B \subset V$ be a subset of the vertices such that $|B|/|V| \leq 1/100$. Let N denote the projection matrix of the set B , and M denote the projection matrix of the set \overline{B} . For $\phi \in \mathbb{R}^{|V|}$ we have

$$(i) |\phi P^t M|_2 \leq |\phi|_2;$$

$$(ii) |\phi P^t N|_2 \leq \frac{1}{3} |\phi|_2.$$

Proof: The eigenvalues of P^t all have absolute value at most 1, so $|\phi P^t|_2 \leq |\phi|_2$. Since M is a projection matrix, for any vector ϕ we have $|\phi M|_2 \leq |\phi|_2$, because multiplying by M effectively just sets some components of ϕ to 0 and does not increase any other components. Combining these two

facts proves (i). To prove (ii), first write $\phi = v + w$ where v is a scalar multiple of the stationary eigenvector $(1, 1, \dots, 1)$ and w is orthogonal to v . Now note

$$|\phi P^t N|_2 = |v P^t N + w P^t N|_2 \leq |v P^t N|_2 + |w P^t N|_2,$$

by linearity and the triangle inequality. Now, working with the right hand side, the first term satisfies

$$|v P^t N|_2 = |v N|_2 \leq \frac{1}{10} |v|_2 \leq \frac{1}{10} |\phi|_2.$$

The first equality is because $v P^t = v$. The second is because $|\overline{B}| \leq \frac{1}{100} |V|$ and v is parallel to $(1, 1, \dots, 1)$. Now to bound the w component, we have

$$|w P^t N|_2 \leq |w P^t|_2 \leq \lambda_*^t |w|_2 \leq \frac{1}{10} |w|_2 \leq \frac{1}{10} |\phi|_2,$$

since by choice of t , we have $\lambda_*^t \leq \frac{1}{10}$. Thus $|\phi P^t N|_2 \leq \frac{1}{5} |\phi|_2$. \blacksquare

Theorem 4.6 (Expander Sample Majorities) [IZ89] *Let P be the strongly aperiodic walk on G_n , started in its uniform stationary distribution π , and let t be as defined in Section 4.1. Let $B \subset V$ be any subset of the vertices such that $|B|/|V| \leq 1/100$. If $X_1, X_2, X_3, \dots, X_{8k}$ are $8k$ samples drawn t steps apart from the walk P , then the probability that at least $4k$ (half) of these samples X_i lie in B is at most 2^{-k} .*

Proof: Notice that, since we start in the uniform stationary distribution π the probability of a getting a given sequence of results in B and out of B , say “Out, In, Out, Out” denoted $B\overline{B}BB$, is given by

$$\Pr\{B\overline{B}BB\} = |\pi P^t M P^t N P^t M P^t M|_1.$$

To bound \mathcal{L}^1 norms, we use the lemma above in combination with the fact that, for $\phi \in \mathbb{R}^{|V|}$ we have $|\phi|_1 \leq \sqrt{|V|} |\phi|_2$, which is obtained by applying the Cauchy-Schwartz inequality.

Call a sequence σ of $8k$ samples ‘bad’ if at least $4k$ of the samples lie in B . We can now bound

$$\Pr\{\text{obtaining a specific ‘bad’ } \sigma\} \leq \sqrt{|V|} |\pi(P^t N)^{4k}|_2 \leq 5^{-4k} |\pi|_2 \sqrt{|V|} = 5^{-4k},$$

when starting the walk in the uniform (stationary) distribution $\pi(v) = \frac{1}{|V|}$ for all $v \in V$.

Because there are at most 2^{8k} sequences of results σ (both ‘bad’ and otherwise), we have

$$\Pr\{\text{obtaining some ‘bad’ } \sigma\} \leq 2^{8k} 5^{-4k} = (256/625)^k \leq 2^{-k},$$

which is the desired result. \blacksquare

This gives us the following corollary by the same reasoning as before.

Corollary 4.7 [IZ89] There is a constant c_3 (specified in Section 4.1), such that there is a deterministic algorithm that given only $n + c_3 k$ uniform independent random bits, outputs $8k$ random binary n -tuples $X_1, X_2, X_3, \dots, X_{8k}$ with the following property.

If B is any subset of $V = \{0, 1\}^n$ such that $|B|/|V| \leq 1/100$, then the probability that at least $4k$ of the samples X_i lie in B is at most 2^{-k} .

Proof: Apply the preceding theorem. The constant c_3 is a bound on the number of random bits required to take $8t$ steps of the walk. \square

Remark: As before, a similar result holds on any similar magnifier. Different constants can be achieved.

The material in this section was adapted from [IZ89]. If B is taken to be the set of incorrect witnesses in a BPP algorithm, the above corollary is precisely the expander-based result in that paper. Equivalent results are found in [CW89]. Both papers use essentially the techniques in [AKS87], but significantly clarifying both the presentation and the importance of the results. \square

Corollary 4.8 Suppose we have an algorithm S to approximate some value a which, when given a word w of s i.i.d. uniform random bits, produces an estimate $S(w)$ of a such that

$$\Pr\{|S(w) - a| > \beta\} \leq 1/100.$$

Then there is an approximation algorithm which, given $s + C(\lg(1/\delta))$ random bits for any real δ , produces an estimate M such that

$$\Pr\{|S(w) - a| > \beta\} \leq \delta.$$

Proof: Let $S(w)$ be the function computed by the algorithm on random input $w \in W = \{0, 1\}^s$. We are guaranteed that when w is chosen uniformly from W that $\Pr\{|S(w) - a| > \beta\} \leq 1/100$. It follows that if we consider the subset $B \subset W$ given by

$$B = \{w \mid |S(w) - a| > \beta\}$$

then $|B|/|W| \leq 1/100$.

So suppose now that we generate $8k$ random elements w_1, w_2, \dots, w_{8k} of W using the method in Corollary 4.7 with the expander G_s . This requires $s + c_3 k$ random bits, and we are guaranteed that the probability that at least half ($4k$) of the samples w_i lie in B is at most 2^{-k} .

It follows that the median M of the $8k$ values $S(w_i)$ must lie in the interval $[h_1 - \beta, h_1 + \beta]$ with probability at least $1 - 2^{-k}$. For, if the median lies outside a given interval, at least half of the $S(w_i)$ do, and so by definition of B , at least half of the w_i are in B . (This is the same reasoning as in Lemma 3.1.) Conclude that if $k \geq \lg(1/\delta)$, the probability that the median M lies within the interval $[h_1 - \beta, h_1 + \beta]$ is at least $1 - \delta$. \square

Remark: This is a constructive version of the BPP result of [IZ89]. \square

4.5 Combining the Results

Combining the results on sample means and sample majorities from expanders gives us the following theorem.

Theorem 4.9 *There are constants c_1 and c_3 (specified in Section 4.1), such that there is a deterministic algorithm that when given any function $h : \{0, 1\}^n \rightarrow \mathbb{R}$ takes*

$$n + c_2 \lceil 100h_2/\beta^2 \rceil + c_3 \lceil \lg(1/\delta) \rceil$$

random bits and produces an estimate M satisfying

$$\Pr\{|M - h_1| > \beta\} \leq \delta.$$

The running time of the algorithm is polynomial in n , $1/\beta$, $\lg(1/\delta)$, h_2 , and the cost of computing h .

Proof: Let $S(w)$ be the estimate computed by the algorithm given by Corollary 4.4, whose input w is the string of $s = n + c_2 \lceil 100h_2/\beta^2 \rceil$ bits. That corollary guarantees us that when w is chosen uniformly at random from $W = \{0, 1\}^s$ that $|S(w) - h_1| > \beta$ with probability at most $1/100$.

The result follows from Corollary 4.8. \blacksquare

4.6 The Implied Algorithm

The reader should already have understood the algorithm implied by the preceding results. However, we specify it in Figure 4.10 for completeness.

The description makes evident a decomposition of the algorithm into two phases: one in which all samples are produced and another in which these samples are used to produce the estimate. This decomposition places the algorithm into a class of approximation algorithms considered in [BGG90, Section 3]. Also it shows that the results here can be considered to provide a pseudo-random generator for the usual median-of-sample-means algorithm. We discuss this in the next section.

4.7 Discussion

Informally, a pseudo-random generator is an algorithm that, given a small uniform random input “seed”, generates a larger (correlated) sequence of random outputs, which are “essentially as good as” true uniform iid values. A cryptographically secure pseudo-random generator would generate outputs which were indistinguishable from true uniform iid values by all polynomial time computations. The existence of such generators is closely linked to the existence of certain “one-way” functions, whose existence, in turn, is a hard and well-known open problem. (See [ILL89].) Therefore

researchers have concentrated on proving that certain pseudo-random generators are good for certain specific applications. To this end, for example, Bach [Bac87] and Karloff and Raghavan [KR88] have shown that certain linear congruential generators are good in specific algorithms, such as taking square roots modulo a prime and QuickSort.

We have shown an algorithm that produces samples in $\{0, 1\}^n$ by taking every i th state of a stationary random walk on G_n , and have demonstrated that this algorithm can be used in a certain precise sense as a pseudo-random generator for computing sample means (see Theorem 4.3).

The authors of [IZ89] (also [CW89]) proved the results of Section 4.4 showing that the same pseudo-random generator could be used in computing certain sample majorities. Their purpose was to show that the generator could be used to rapidly decrease the error probabilities of any BPP algorithm. These techniques do not seem able to give Theorem 4.3 directly.

One can use either the results of Section 4.3 or of Section 4.4 alone with the median-of-sample-means algorithm. But neither does as well as the combination. Applying Theorem 4.3 alone with the median-of-sample-means algorithm, one can estimate M within β of h_1 with probability $1 - \delta$ using $\Theta((n + h_2/\beta^2) \lg(1/\delta))$ independent random bits. Similarly, applying Corollary 4.7 alone with the median-of-sample-means algorithm one can get away with using $n + O((h_2/\beta^2) \lg(1/\delta))$ independent random bits, but not fewer. However we have shown that the same generator (with different sizes of G_n) can be used twice in succession to yield an extra savings so that only $n + O(h_2/\beta^2) + O(\lg(1/\delta))$ bits are used.

We got this last idea from the recent paper [BGG90], in which the authors show how to combine the expander-based pseudo-random generator presented here with certain pairwise-independence constructions. Their construction gets even better results. For boolean functions $h : \{0, 1\}^n \rightarrow \{0, 1\}$ they give a pseudo-random-generator for the same median-of-sample-means algorithm that produces $\Theta((1/\beta^2) \lg(1/\delta))$ samples using $2n + O(\lg(1/\delta))$ random bits. This works because they show that using only $2n$ bits, they can generate samples whose sample mean is within β of h_1 with probability bounded away from zero. Then by Corollary 4.8, the result follows.

Notice that their random bit requirement is independent of β . They assume (critically) that $\beta \geq 2^{-n/2}$ in order that the pairwise independence construction be able to produce enough samples. This is a reasonable assumption, since otherwise the number of samples produced by this method would be comparable to the size of the sample space, the total running time of the method would be exponential, and it would not make sense to use this method over a simple exhaustive exact computation of h_1 .

An examination of their proof reveals that their method also works with arbitrary functions $h : \{0, 1\}^n \rightarrow \mathbb{R}$, under the slightly stronger assumption that $\beta^2/h_2 \geq 2^{-n}$, which is reasonable for the same reason. (Indeed, it is reasonable to assume that h_2/β^2 is bounded by a polynomial.)

Any approximation scheme of the form: (I) produce several samples (II) estimate mean using the samples, which works for every boolean function h in polynomial time, must produce $\Omega(\beta^2 \lg(1/\delta))$

samples. Every such scheme that produces this many sample points requires $\Omega(n + \lg(1/\delta))$ random bits [BGG90]. Thus one cannot do better in the same framework.

Algorithm 4.10 Random-Bit Efficient Estimation The setup described in Section 4.1 is assumed. We assume also that we have an oracle computing the function h . This algorithm takes

$$n + c_2 \lceil 100h_2/\beta^2 \rceil + c_3 \lceil \lg(1/\delta) \rceil$$

random bits and outputs an estimate M of the mean value h_1 of h satisfying $\Pr\{|M - h_1| > \beta\} \leq \delta$.

Phase I. (Produce Samples)

Stage 1. (Produce Sampling Seeds)

- Use $s = n + c_2 \lceil 100h_2/\beta^2 \rceil$ input random bits to specify an initial vertex of a random walk on G_n . The vertex set is $W = \{0, 1\}^n$. This starts us in the uniform stationary distribution of the walk. The $c_3 \lceil \lg(1/\delta) \rceil$ additional input random bits that we have are sufficient to take all of the steps of the following walk.
- for $1 \leq i \leq 8 \lceil \lg(1/\delta) \rceil$ do begin
 - Take t steps of the walk on G_n . After the t th step, let (sampling seed) w_i be the current vertex, a string of s bits.

Stage 2. (Produce Samples from Sampling Seeds)

- for $1 \leq i \leq 8 \lceil \lg(1/\delta) \rceil$ do
 - Use the first n bits of w_i to specify an initial vertex of a random walk on G_n whose vertex set is $V = \{0, 1\}^n$. This starts us in the uniform stationary distribution of the walk. The $s - n = \lceil 100c_2h_2/\beta^2 \rceil$ remaining bits of w_i are sufficient to take all of the steps of the following walk.
 - for $1 \leq j \leq \lceil 100h_2/\beta^2 \rceil$ do
 - Take t steps of the walk on G_n . After the t th step, let the sample X_{ij} be the current vertex, a string of n bits.

end

Phase II. (Produce Estimate from Samples)

- for $1 \leq i \leq 8 \lceil \lg(1/\delta) \rceil$ do
 - Let $\alpha_i = \left(\sum_{1 \leq j \leq \lceil 100h_2/\beta^2 \rceil} X_{ij} \right) / \lceil 100h_2/\beta^2 \rceil$ be the sample mean of the X_{ij} for the given i .
 - Let M be the median of the α_i . Return the estimate M .

Figure 4.10: Random-Bit Efficient Estimation

Chapter 5

Estimating the Significance of Contingency Tables

A two-way contingency table is a tabular description of the sizes of intersections of two partitions of a set of N elements. Given two such partitions, $\mathcal{R} = R_1, R_2, \dots, R_m$ and $\mathcal{C} = C_1, C_2, \dots, C_n$, we get an $m \times n$ table T by letting

$$T_{ij} = |R_i \cap C_j|.$$

These tables arise in statistical situations where one wants to investigate possible correlations between two partitions of interest. For example, the object of a certain medical study may be to investigate relations between smoking and the incidence of heart attacks. A sample is taken from a set of subjects. The subjects are classified as non-smokers, light smokers, and heavy smokers by the number of cigarettes consumed daily. The subjects are also classified as having had 0, 1, or more than 1 heart attack. This may result in a table like that in Figure 5.1. (Note: the data in that table are entirely fictitious.)

smoking	heart attacks		
	0	1	more
none	621	3	16
light	84	25	3
heavy	64	37	9

Figure 5.1: a 3×3 contingency table

Viewing the table, there is an apparent correlation between heavier smoking and heart attacks. But is this indeed the case? Might we expect to see a similar table even if the two were independent? One way of testing for independence is to perform a statistical test using a measure of independence

called the chi-squared (χ^2) statistic.

Let $r_i = |R_i| = \sum_j T_{ij}$ and $c_j = |C_j| = \sum_i T_{ij}$. Suppose we adopt a so-called null hypothesis that the table T had been produced simply by sampling subjects with replacement from a population with a fraction r_i/N falling in category R_i and a fraction c_j/N falling into category C_j , with the two classifications being completely independent. This is called the multinomial model. (For additional background see [BFH75, Eve77, MGB74].) Then having observed the table T , the maximum likelihood estimate for $|R_i \cap C_j|/N$ is $\bar{T}_{ij} = r_i c_j / N$. This is also the expected value of T_{ij} if it were generated at random under the multinomial model. The statistic $\chi^2(T)$ measures the distance of the observed table T from this outcome; it is defined

$$\chi^2(T) = \sum_{i,j} \frac{(T_{ij} - \bar{T}_{ij})^2}{\bar{T}_{ij}}.$$

We define the multinomial significance of the observed table T as $\Pr\{\chi^2(T') \leq \chi^2(T)\}$ when T' is drawn under the hypothesized multinomial model. The greater this value, the more *unlikely* it is that we would see a table as "skewed" as T , had the table simply been generated at random under the multinomial model. We reject the hypothesis that the table was produced at random under this model if the significance is larger than a certain threshold, typically .90 or .95. We might then conclude that T in fact shows some relation between the partitions \mathcal{R} and \mathcal{C} . This is called the exact multinomial χ^2 test. For large N , under the multinomial model, the distribution of $\chi^2(T)$ approximately follows a distribution called 'the χ^2 distribution with $(m-1)(n-1)$ degrees of freedom,' [MGB74]. The standard χ^2 test uses this approximation to estimate the multinomial significance.

Recently there has been interest in testing for significance against a different null hypothesis, one in which the underlying model is uniform. In this model all tables with the same row and column sums arise with equal probability. Again, the significance value of the observed table T is $\Pr\{\chi^2(T') \leq \chi^2(T)\}$, but here drawing T' uniformly from the set of all tables with the same row and column sums. Again, we reject this null hypothesis if the significance value is large; this test is called a conditional volume test. (Diaconis and Efron [DE85] give several reasons why this test is meaningful. The curious reader is referred to that article. The relevant statistics goes well outside the scope of this work.)

Statisticians also employ tests based on the significance in other underlying models. Several algorithms have been suggested for estimating and exactly calculating significance levels for the χ^2 test under the multinomial, the Fisher-Yates (hypergeometric), and other models (e.g., see [PTH81, MPS88, BOP88]). It seems significantly harder to compute significance values under the uniform model. There do not seem to be any non-exhaustive methods known to compute the exact significance. Moreover, it seems that one should not seek a randomized method that relies on exact uniform sampling. We say this because computing the number of tables with the same row and column sums as T is a well-known and difficult combinatorial problem and most traditional

methods of exact sampling would rely on methods of computing this number exactly. (Some combinatorial background material appears in Appendix C. Two good references are [GC77] and [Sta86, Chapter 4, esp. p. 232]. A provable relationship between the complexity of sampling and approximate counting in general is described in [JV86]. We explain the basic idea behind approximate counting using sampling in Appendix C.)

In this chapter we present a randomized algorithm for estimating the significance of tables under the uniform model. Our method is based on sampling using a random walk. We can prove that the walk converges to the desired uniform distribution, so that asymptotically the method must yield good estimates, but we cannot currently give good (polynomial) time bounds on the rate of convergence. However, we conjecture some bounds and show that empirically the method seems to perform well under the conjectured bounds.

Let us now fix some notation. Let $\mathbf{r} = (r_1, r_2, \dots, r_m)$ and $\mathbf{c} = (c_1, c_2, \dots, c_n)$ denote nonnegative integer partitions of N , with $r_1 \geq r_2 \geq \dots \geq r_m > 0$ and $c_1 \geq c_2 \geq \dots \geq c_n > 0$. Let $\Sigma_{\mathbf{r}\mathbf{c}}$ denote the set of all $m \times n$ nonnegative integer matrices in which row i has sum r_i and column j has sum c_j . (We have permuted the rows and columns so that the sums are in non-increasing order. The properties of the tables that concern us, the cardinality $|\Sigma_{\mathbf{r}\mathbf{c}}|$ and the χ^2 statistic, do not depend on the order of the rows and columns.) Note that we have $\sum_{i,j} T_{ij} = N$ for every $T \in \Sigma_{\mathbf{r}\mathbf{c}}$, and also that m and n , the dimensions of the tables, are not explicitly present in the notation $\Sigma_{\mathbf{r}\mathbf{c}}$ but are understood in context from the lengths of \mathbf{r} and \mathbf{c} . This use of m , n , and N will persist throughout the chapter.

The set $\Sigma_{\mathbf{r}\mathbf{c}}$ is always nonempty. This may not be evident immediately, but an algorithm for constructing an element is described briefly in Section C.1.4 of Appendix C. We will also assume that $m > 1$ and $n > 1$; otherwise $\Sigma_{\mathbf{r}\mathbf{c}}$ has only one element.

In Section 5.1, we present a random walk on the set $\Sigma_{\mathbf{r}\mathbf{c}}$. We can prove that the walk converges to the uniform distribution on $\Sigma_{\mathbf{r}\mathbf{c}}$. However, we cannot currently prove polynomial time bounds for the time to reach stationarity.

In Section 5.4, we present the results of some empirical studies involving the algorithm. We compare results obtained by our randomised method with those obtained by exact methods, (when they are feasible). We also compare our results with asymptotic approximations obtained by some analytic methods.

In Section 5.5 we discuss how recent results by some other researchers may be applied to prove polynomial time bounds for a variant of our scheme. The polynomials in the bounds, however, are too large to offer both rigorous and practical utility.

5.1 A Random Walk on Σ_{rc}

In this section we suggest a random walk on Σ_{rc} that can be used for near-uniform sampling. We prove that the walk converges to the uniform distribution on Σ_{rc} , but we do not prove a polynomial convergence rate bound for this walk.

The walk can start anywhere in Σ_{rc} . In practice one has 'observed' a table in Σ_{rc} as the result of some statistical study, and this can be used as an initial element at which to start the walk. Alternatively, it is not hard to construct a table by the algorithm described in Appendix C, Section C.1.4. After an initial element is obtained, each step of the walk is generated by performing the following operations.

Algorithm 5.2 [Basic random walk on Σ_{rc}]. Procedure to take a single step of the walk on Σ_{rc} .

1. Let $X \in \Sigma_{rc}$ be the current position of the walk.
2. Choose a pair of rows i_1 and i_2 , with $1 \leq i_1 < i_2 \leq m$, uniformly from among all $\binom{m}{2}$ such choices.
3. Choose a pair of columns j_1 and j_2 , with $1 \leq j_1 < j_2 \leq n$, uniformly from among all $\binom{n}{2}$ such choices.
4. Choose d uniformly from the set $+1, -1$.
5. Let Y be the $m \times n$ matrix obtained from X by taking

$$Y_{i_1, j_1} = X_{i_1, j_1} + d \quad Y_{i_1, j_2} = X_{i_1, j_2} - d$$

$$Y_{i_2, j_1} = X_{i_2, j_1} - d \quad Y_{i_2, j_2} = X_{i_2, j_2} + d$$

and $Y_{i,j} = X_{i,j}$ for all other $(i,j) \in [m] \times [n]$. Notice that Y has the same row and column sums as X , and is a matrix of integers.

6. (Take the Step) If all entries of Y are nonnegative, then move to Y ; our new current position is Y . Otherwise remain at X , i.e. our new current position remains X . Note: The entries of Y will in fact be nonnegative, unless X has a zero in at least one of the two entries from which we are subtracting 1.

Since each of the possible moves preserves the row and column sums of the matrix, (adding and subtracting 1), this walk always keeps us in the set Σ_{rc} . We now show that these steps allow us to move between any two elements in Σ_{rc} .

Lemma 5.3 *The walk whose steps are generated by the procedure above is irreducible on Σ_{rc} . Moreover, if (i,j) is the lexicographically first coordinate in which A and B differ, then there is a path between A and B that does not alter coordinates lexicographically preceding (i,j) .*

Proof: We need to show that for each X and Y in Σ_{rc} , there is a path between X and Y , using only possible steps of the walk. Note that since each step is reversible, a path from X to Y implies a symmetric one from Y to X .

We prove that there is a path joining X and Y by induction on a distance measure between the two tables X and Y . Define the distance $d(X, Y)$ between X and Y as $d(X, Y) = \sum_{i,j} |X_{i,j} - Y_{i,j}|$. Observe that $d(X, Y) = 0$ if and only if $X = Y$. Further note that, since the grand sums in the table are the same, this distance is always a multiple of two.

Let k be a nonnegative integer, and assume the following induction hypothesis: if $0 < d(X, Y) \leq 2k$ and if (i, j) is the lexicographically-first coordinate in which X and Y differ, then there is a path joining X and Y using only steps of the walk that do not involve coordinates lexicographically preceding (i, j) . This is vacuously true for $k = 0$.

For the induction step, let X and Y be two elements of Σ_{rc} and suppose $d(X, Y) = 2(k + 1)$, where (i, j) is the first coordinate in which they differ. We will show that either there is a move from X to X' where $d(X', Y) \leq 2k$ or there is a move from Y to Y' where $d(X, Y') \leq 2k$, where no coordinates preceding (i, j) are involved. The induction hypothesis will then imply that there is an entire path between X and Y .

In step 2 of the algorithm choose $i_1 = i$ and $j_1 = j$. Then there are two cases to consider.

Case (a) $X_{i_1, j_1} < Y_{i_1, j_1}$. Then since each row and column of X has the same sum as in Y , we have

$$\exists j_2 \text{ such that } X_{i_1, j_2} > Y_{i_1, j_2}$$

$$\exists i_2 \text{ such that } X_{i_2, j_1} > Y_{i_2, j_1}$$

We must have $i_1 < i_2$ and $j_1 < j_2$, since (i_1, j_1) was chosen as the lexicographically first position in which X and Y differ. Moreover, the entries X_{i_1, j_1} and X_{i_2, j_1} are both positive, since they are greater than their nonnegative counterparts in Y . This means, that letting $d = +1$, the move

$$X'_{i_1, j_1} = X_{i_1, j_1} + 1 \quad X'_{i_1, j_2} = X_{i_1, j_2} - 1$$

$$X'_{i_2, j_1} = X_{i_2, j_1} - 1 \quad X'_{i_2, j_2} = X_{i_2, j_2} + 1$$

yields an X' having nonnegative entries as well as sharing the same row and column sums as X .

By moving from X to X' , the difference with respect to Y on least the three coordinates (i_1, j_1) , (i_1, j_2) , and (i_2, j_1) decreased by 1. The difference at (i_2, j_2) may have increased by 1, but the net change in all four coordinates must in any case be a decrease of at least 2. That is, $d(X', Y) \leq d(X, Y) - 2$, so $d(X', Y) \leq 2k$. Now by the induction hypothesis there is a path from X' to Y . Adding the step from X to X' completes the path from X to Y , without altering any coordinates lexicographically preceding (i, j) (in which X and Y already agree).

Case (b) $X_{i_1, j_1} > Y_{i_1, j_1}$. This case is entirely symmetric. Swapping the roles of X and Y , the same argument as in case (a) shows that there is a move from Y to Y' with $d(X, Y') \leq 2k$. Thus,

by the induction hypothesis, there is a path between X and Y' , and hence a path between X and Y via Y' . \blacksquare

Combining the above lemma with the fact that the Markov chain is symmetric yields the following desirable result.

Theorem 5.4 *The walk whose steps are generated by the procedure of Algorithm 5.2 is ergodic and has uniform stationary distribution on Σ_{rc} .*

Proof: If P_{XY} is the probability of moving from X to Y in one step, then $P_{XY} = P_{YX}$, since if we are at Y , then to move to X we need to pick the same rows and columns as when moving from X to Y , but pick d with the opposite sign. Thus the walk is a symmetric Markov chain, and its unique stationary distribution will therefore be uniform, provided that the walk is in fact ergodic. (See Chapter 1.)

In order to show that the walk is ergodic, we need to show it is irreducible and aperiodic. We showed the walk was irreducible in Lemma 5.3. The walk cannot be periodic, since there are necessarily some states in which there is a possibility of immediate return. To see this notice that starting at any table and making any given move repeatedly, one can proceed at most N times before one of the two entries from which we are repeatedly subtracting 1 reaches zero. From such a state, choosing the same move will cause us to remain at the same position for that step, an immediate return. Hence that state is aperiodic. Since the chain is irreducible, the aperiodicity of one state implies that every state is aperiodic. \blacksquare

From this and the Basic Convergence Theorem (1.1), we can conclude that if we run the walk long enough, we can certainly use it to obtain uniform samples.

5.2 Bigger Steps

The basic walk uses unit step sizes. That is, each step of the walk only adds or subtracts 1 from the entries it affects. Intuition tells us that allowing moves of step sizes greater than 1 will increase the rate of convergence. But one must be somewhat careful in 'designing' a walk with larger possible step sizes. We wish that our walk be time-reversible (so that we can apply our earlier results) and have uniform stationary distribution. This will be the case if and only if the chain is symmetric. The choice of the step size should be such that the probability of a single step of the walk from a table T_1 to a table T_2 and the step's reversal from T_2 to T_1 are the same. A naive method may disrupt the symmetry of the walk and make the stationary distribution non-uniform. For example, one might initially think of the following strategy: At each step choose a random step size between 1 and the minimum entry on the four corners of the move. This walk will remain aperiodic and

irreducible, but it will no longer have symmetry. The probability of a step and its reversal will not always be the same.

One can legitimately choose the step size based on the sum (or average) of the entries on the four chosen corners of the move; this sum remains invariant under the move, and this guarantees that the move and its reversal obey the required symmetry. In our experiments, we actually use the following walk with random step sizes.

Algorithm 5.5 [Random walk with larger step sizes]. Procedure to take a single step of the walk on Σ_{rc} using random step sizes.

1. Let the $X \in \Sigma_{rc}$ be the current position of the walk.
2. Choose a pair of rows i_1 and i_2 , with $1 \leq i_1 < i_2 \leq m$, uniformly from among all $\binom{m}{2}$ such choices.
3. Choose a pair of columns j_1 and j_2 , with $1 \leq j_1 < j_2 \leq n$, uniformly from among all $\binom{n}{2}$ such choices.
4. Choose d uniformly from the set $+1, -1$
5. Let $a = \lceil X_{i_1, j_1} + X_{i_1, j_2} + X_{i_2, j_1} + X_{i_2, j_2} / 32 \rceil$, and choose s uniformly from $\{1, \dots, a\}$.
6. Let Y be the $m \times n$ matrix obtained from X by taking

$$Y_{i_1, j_1} = X_{i_1, j_1} + ds \quad Y_{i_1, j_2} = X_{i_1, j_2} - ds$$

$$Y_{i_2, j_1} = X_{i_2, j_1} - ds \quad Y_{i_2, j_2} = X_{i_2, j_2} + ds$$

and $Y_{i, j} = X_{i, j}$ for all other $(i, j) \in [m] \times [n]$. Notice that Y has the same row and column sums as X , and is a matrix of integers.

7. (Take the Step) If all entries of Y are nonnegative, then move to Y ; our new current position is Y . Otherwise remain at X , i.e. our new current position remains X . Note: The entries of Y will in fact be nonnegative, unless $X < s$ in at least one of the two entries from which we are subtracting s .

An argument similar to the one in Theorem 5.4 shows that the modified algorithm also gives an ergodic walk with uniform stationary distribution on Σ_{rc} .

The constant 32 that is used in step 5 is somewhat arbitrary. We chose it by experimentation from among a small set of values tested in walks on some typical sets Σ_{rc} that we wished to study; the value 32 was small enough that in step 5 we did not simply get $a = 1$ most of the time, and 32 was large enough that a high ratio of moves attempted in step 7 were actually taken. In general, these properties depend on the range of values appearing in the tables.

5.3 Eigenvalue Hypothesis

We are currently unable to prove general polynomial upper bounds on the convergence rate of the walk. So we are not in a position to say precisely how long we must run the walk to obtain good samples. We can prove bounds in some special cases and for some analogous walks, and we discuss these later. In this section, we motivate an hypothesis about the second absolute eigenvalue of the walk on Σ_{rc} . Under the hypothesis we can calculate bounds on the convergence rate and the number of samples needed to obtain given variance in our sample means. We verify empirically that these seem to be valid, though not necessarily tight.

For any table $T \in \Sigma_{rc}$, knowing the $d = (m-1)(n-1)$ entries T_{ij} , with $i \neq 1$ and $j \neq 1$ allows us to determine the remaining entries by the sum constraints. We can thus specify the table T as a point of the integer lattice in d -dimensional space, and the set Σ_{rc} can be associated to the set of lattice points in the d -dimensional convex polytope given by the sum constraints. Each entry T_{ij} must satisfy $0 \leq T_{ij} \leq \min(r_i, c_j)$, so it follows that the entire convex region sits within the corresponding d -dimensional "bounding box". Our walk(s) on Σ_{rc} are walks on these lattice points, where the set of possible transitions are a superset of the set of lattice edges. (A lattice edge joins two integer lattice points whose distance is 1.)

Conjecture 5.6 (Eigenvalue Hypothesis) *For the walk on Σ_{rc} described in Algorithm 5.5, the second absolute eigenvalue λ_2 satisfies*

$$\frac{1}{1 - \lambda_2} \leq \frac{1}{9} \sum_{1 < i \leq m, 1 < j \leq n} \min(r_i, c_j)^2.$$

Our conjecture is motivated by two intuitive lines of thought. First, one expects that the conductance (see Chapter 1) of subgraphs of the lattice grid that are delimited by convex polytopes will have conductance matching that of the grid. The walk on the d -dimensional lattice 'box' whose sides are $\min(r_i, c_j)$ ($1 \leq i \leq m, 1 \leq j \leq n$) would satisfy the conjecture. For example, in the case of a 2×2 table with row sums r_1, r_2 and column sums c_1, c_2 , the basic walk of Algorithm 5.2 may be viewed as the natural random walk on a line segment of length $\ell = \min\{r_1, c_1\} - \max\{0, r_2 - c_1, c_2 - r_1\}$.

Our second motivation comes from an analogy with the limiting, continuous case. Payne and Weinberger [PW60] proved the following theorem.

Theorem 5.7 (Payne-Weinberger) *Let D be a convex body with diameter σ in \mathbb{R}^n . Then*

$$\inf \frac{\int_D |\nabla f|^2 dD}{\int_D f^2 dD} \geq \frac{\pi^2}{\sigma^2}, \quad (5.1)$$

where the infimum is over all functions f with bounded second-derivative on D , and satisfying $\int_D f dD = 0$, and here ∇f denotes the real gradient operator and $\|\cdot\|$ is the L^2 norm on \mathbb{R}^n .

The theorem applies to an arbitrary convex body. When it so happens that the boundary of D is sufficiently smooth, the infimum in 5.1 gives the minimax characterization of the eigenvalue of

the Laplacian operator ($\sum_{i=1}^n \frac{\partial^2}{\partial x_i^2}$) that is smallest over all functions whose gradient field exhibits no flow across the boundary of D . (Commonly called a Neumann condition, this is expressed by requiring $\frac{\partial f}{\partial n} = 0$ everywhere on the boundary of D , where at each point on the boundary n represents the outward normal vector.) However, to avoid smoothness requirements, one may work strictly with 5.1, which holds for all functions satisfying the stated conditions, whether or not this infimum corresponds to an eigenvalue of ∇^2 on D .

The reader will note the similarity of 5.1 to 1.8. It is a direct analogue. Indeed, the Q -matrix of a finely-spaced lattice graph (mesh) within a body D is commonly-used difference approximation to the Laplacian operator on the D , (see [BL84]). Fine meshes correspond, by scaling, directly to the case when N is large. It is, therefore, reasonable to believe that an analogous bound holds for lattice graphs embedded within large convex bodies. In Section 5.6 we describe an approximation conjecture that would imply the eigenvalue bound.

Dyer, Frieze, and Kannan [DFK89] proved a bound on the eigenvalues of walks on lattice graphs within certain convex sets. Their bound requires that the region be "well-rounded" (in a certain precise sense) and is somewhat weaker, but is still useful to us. We discuss this in Section 5.5.

Remark: The author tried unsuccessfully to show a specific eigenvalue bound for the basic walk on Σ_{rc} , using the 'canonical path' methods presented in Chapter 1 on the paths described in Theorem 5.4. \square

5.4 Experimental Results

The results in this section are experiments using computer simulations of the random walk. The reader should find that they lend supportive evidence to the eigenvalue hypothesis, and demonstrate the practicality of the algorithms suggested. All of our experiments seem to perform at least as well as we should expect using our bounds, and the conjectured eigenvalue bound. We ran several more experiments than those we present here. However, most of the relevant issues seem to be covered by this set of examples.

These experiments were all conducted on a DECstation 3100, a desktop workstation that executes about $(1.2 \pm .2) \times 10^7$ VAX-equivalent machine instructions per second. The experiments were run in a time-sharing environment. This means that the process computing the experiment shares CPU time with other processes on the system, including other user's tasks. The CPU figures presented with each experiment are the 'user' CPU seconds actually used by the process. (This excludes CPU time used by kernel-level tasks initiated by the process). The effect of timesharing in the CPU time figures is negligible. Real time figures are also presented; these, however, are quite dependent upon the amount of system load incurred by other processes running concurrently, but since in most cases the experiment was allotted more than 90% of the CPU cycles that elapsed during its execution, this effect is also small.

Random numbers were generated using a non-linear feedback generator provided by the system. The period of the generator is very large, approximately 3.4×10^{10} , and low-order portions of the generated values are also purported to look random. The time of day was used as a 'random' initial seed.

The programs for the experiments were coded in 'C.' We wrote code for Mathematica (TM) to compute Diaconis and Efron's analytical approximations. These may be obtained from the author.

5.4.1 Background

In this discussion, we assume background material from Chapters 1 and Chapter 3. We also refer to the Kolmogorov-Smirnov distance between two distributions. This is defined and discussed in Appendix A.

Fix Σ_{rc} and an 'observed' table $T \in \Sigma_{rc}$. Let λ_1 be the second-largest eigenvalue of the Markov chain whose steps are described by Algorithm 5.5, and let $\tau = 1/(1 - \lambda_1)$.

Let $H = \{T' \mid \chi^2(T') \leq \chi^2(T)\}$ be the subset of Σ_{rc} consisting of tables with smaller χ^2 values. If h is the indicator function of H , then the significance of T is

$$p = h_1 = |H|/|\Sigma_{rc}|,$$

the mean value of $h(T')$ when T' is chosen uniformly from Σ_{rc} . The results of Chapter 3 and Section 3.3 tell us that for any $\beta > 0$ we can get an estimate M satisfying

$$\Pr\{|M - p| > \beta\} > .96 \quad (5.2)$$

by letting M be the median of 5 independent sample means each based on $4(2\tau - 1)p(1 - p)\beta^{-2}$, which is at most $(2\tau - 1)\beta^{-2}$ adjacent samples drawn from the chain in the stationary distribution.

The value of $\frac{1}{1 - \lambda_1}$, and hence also τ , is bounded by the eigenvalue hypothesis. To calculate the number of steps to use to get near stationarity, we multiplied our conjectured bound on $\frac{1}{1 - \lambda_1}$ by $\ln(\prod_{i,j} (1 + \min(r_i, c_j)))$ which is a bound on $\ln|\Sigma_{rc}|$, (see Theorem 1.8).

The randomized experiments we describe here are all based on the following scheme. The random walk of Algorithm 5.5 is simulated for a number of steps determined by the aforementioned formula to get near stationarity. Then some multiple of 2τ samples are drawn by taking adjacent states of the chain. These samples are used in two ways: (a) the precise proportion of samples whose χ^2 value is less than that of the input table is recorded, (b) a bin is associated with each interval of width 0.5 of the real line, and the number of samples whose χ^2 values fall in each bin is counted to make a sample histogram. This is repeated five times. The significance estimate we report is the median of the sample means recorded in (a), and accuracy bounds are based on (5.2). The histogram we speak about is obtained by taking the median count in each bin of the five histograms. Computing times that we report include the time consumed in all five sample runs, computation of the median significance value, and median histogram.

Where exact distributions and counts were obtained, the underlying algorithm is the exhaustive enumeration method described in [PTH81]. (They also describe a non-exhaustive method for determining exact multinomial significance.) The exhaustive method was used similarly to produce two values: (a) an exact significance value, being the proportion of tables whose χ^2 value was smaller than that of the input table (b) and an exact histogram, which counts the number of tables whose χ^2 values fall in each bin of width 0.5.

The reader should note that the histograms are subject to a discretization error introduced by taking bins of width 0.5, and this will have an effect on the Kolmogorov-Smirnov distances we report. The significance values, both exact and estimated, on the other hand are based on specific counts of the χ^2 values less than that of the input table, and are not subject to the same errors. The inherent error in these values due to limited-precision floating-point arithmetic on the computer should be negligible.

5.4.2 Pinckney Gag Rule

The table in Figure 5.8 was taken from [BFH75, p. 99]. It records a vote taken in 1836 in the U.S. House of Representatives on the "Pinckney Gag Rule", which would have banned certain anti-slavery petitions. States were classified as northern, border, and southern states. By looking at this vote can one detect the presence of a North-South bias on this issue roughly 25 years before the Civil War?

	Yea	Abstain	Nay	r_i
North	61	12	60	133
Border	17	6	1	24
South	39	22	7	68
c_j	117	40	68	$N = 225$

Figure 5.8: 1836 Vote on 'Pinckney Gag Rule'

There is apparent percentage-wise diagonal trend running from South-Yea to North-Nay. The χ^2 statistic is invariant under permutations of the rows and columns, thus one cannot hope to use it to validate the observed diagonal trend. However, one can test whether the amount of 'imbalance' in the table is likely under models in which the rows and columns are independent.

The observed table has $\chi^2(T) \approx 41.08$ which places it higher than the 95th percentile of the χ^2 distribution on 4 degrees of freedom. Its significance in the standard χ^2 test is therefore greater than .95.

The conditional volume test, however, accords it substantially less significance. An exhaustive enumeration of the tables is feasible here, and was computed in 73.7 user CPU seconds and 75

seconds of real time. There were 545025 tables with the same row and column sums as the observed table. The observed table has significance .2959. That is, a little less than 30% of the tables with the same row and column sums have χ^2 values no larger than that of the observed table.

We used the random walk method with 113,017 steps for each re-randomization and $2\tau = 6584$ samples for each of the five contributing histograms. The entire process took 28.2 user CPU seconds and 34 seconds of real time, and yielded the median histogram of Figure 5.9.

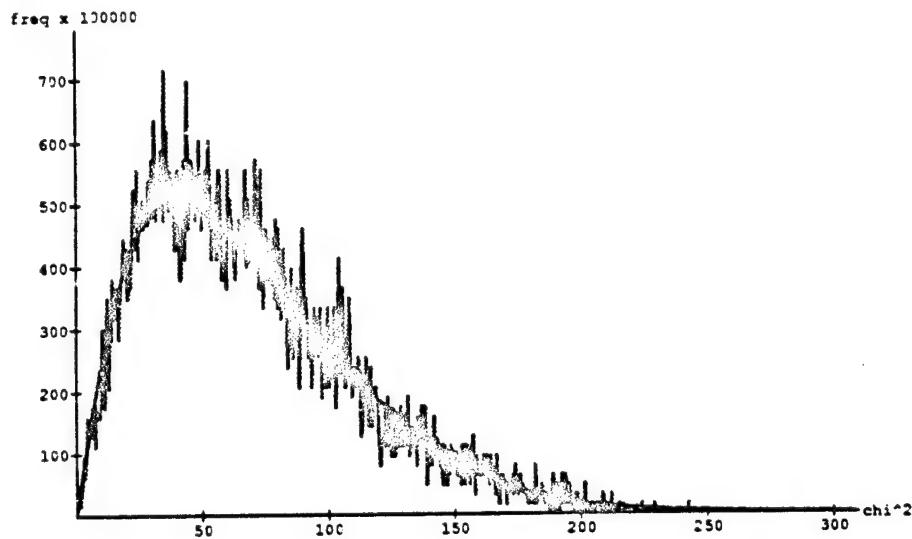


Figure 5.9: Exact (dark) and sample (lighter) distributions of χ^2 under the uniform distribution on Σ_{rc} with margins corresponding to the 'Pinckney Gag Rule' table shown in Figure 5.8. The densities were plotted by joining the histogram data points. Both histograms are normalized here so that each has total mass 100000. A difference of 200 in the y-coordinate marks a difference of only .002 in density. Kolmogorov-Smirnov distance between the two histograms is 0.021. The table has $\chi^2 = 41.08$ which falls at the 29.6% significance level in the exact distribution and 29.8% in the sample distribution.

Our bounds give essentially no usable guarantee on the accuracy of the histogram. They tell us the significance will be within 1 of the true significance with probability at least $31/32$. But the significance takes on values between 0 and 1, so this bound is trivial.

In actuality, however, the Kolmogorov-Smirnov distance between the two histograms in Figure 5.9 is only 0.021. This means the difference in the significance based on the exact and sample histograms is at most 0.02 for any table T . In fact, we find that our estimate for the observed table T is .2982, which agrees to two significant figures with the exact significance.

In order to have guaranteed performance this good, namely an error of at most .01 under the

hypothesized eigenvalue bound knowing that the true value of $p = .30$, we would expect to require $8400 \approx 4 * p(1 - p) * (1/.01)^2$ times as many samples in each run. In this situation, our bounds can only show that four runs of 55,305,600 samples each would be sufficient to guarantee that the estimate is within .01 of the true value with probability greater than .96. This would take about 3 days of real time on the same computer, so it would far better to calculate the distribution exactly for this case. The performance we get here, however, seems to be far better than anything we have proved, even under the hypothesized eigenvalue bound.

5.4.3 Hair Color v. Eye Color

The table shown in Figure 5.10 is discussed in [DE85]. It records the hair color and eye color of 592 subjects. One expects a significant correlation. The χ^2 value of the table is 138.29 which places its multinomial significance higher than .90.

Eye Color	Hair Color				r_i
	Black	Brunette	Red	Blond	
Brown	68	119	26	7	220
Blue	20	84	17	94	215
Hazel	15	54	14	10	93
Green	5	29	14	16	64
c_j	108	286	71	127	N=592

Figure 5.10: Hair color and eye color of 592 subjects

Using an approximation formula in [DE85], (see Appendix C, Section C.1.3) we estimate that the number of tables with matching margins is greater than 10^{15} . This is too large to enumerate exhaustively, so that an exact distribution is not available for comparison. Diaconis and Efron [DE85] obtain an analytic estimate of .41. They know this to be an overestimate. They reduce this to .25 by estimating the effect of what they call 'protrusion'. Using their same formula, we re-calculated their protrusion estimates, and found the reported protrusion estimates to be in error. Using our revised protrusion estimate, the table has significance between .16 and .41. We believe that .16 is pretty close to the truth.

Diaconis and Efron also believed their .25 figure to be an over-estimate. To try to get a handle on the actual value, they used a different randomized algorithm which gives an estimate of .09. Their method produces tables from Σ_{rc} according to the Fisher-Yates (hypergeometric) distribution, with an analytic approximate correction to the uniform. They were able to run their algorithm for smaller values of N ($N = 40$, $N = 60$, and $N = 80$), scaling the table entries proportionally. Then they extrapolated to the actual value of $N = 592$. In our section on scaling, below, we discuss why we believe this was an under-estimate.

Our new method estimates the significance at .15. The median histogram is shown in Figure 5.11. This was computed using $T = 1799772$ to randomize and $2\tau = 46,148$ samples in each of the five contributing histograms. The entire process took 457.16 user CPU seconds and ran in just under 8 minutes of real time. Assuming that the true value is indeed close to .15, our bounds would require 5100 times as many samples, or about 2.35×10^8 samples per run, in order to guarantee an error of at most 0.01 with probability greater than .96. This would take an estimated 28 days on the same computer (though achieving the accuracy 0.01 with confidence about .75 would only take about 5.6 days). Here are two points for comparison: (1) the statistical studies from which such tables arise often take several months; (2) an exhaustive enumeration on the same computer at an estimated rate of 500,000 tables per minute would take about 3805 years. A factor of 5 to 10 speedup in the computational power of desktop computing over the next decade would make the random walk method much more reasonable. Any method that relies on enumerating a substantial fraction of the tables, however, would remain equally unreasonable.

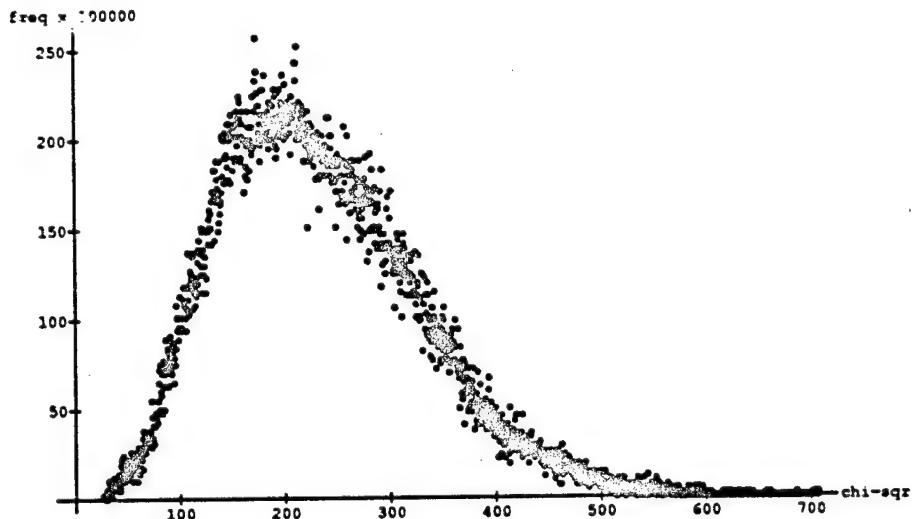


Figure 5.11: Sample histogram of χ^2 under the uniform distribution on Σ_{rc} with margins as in Figure 5.10. Total mass is again normalized to 100000. The smoothness and essential completeness of the curve suggests that the walk sampled a representative selection of the approximately 10^{15} tables in Σ_{rc} .

5.4.4 Irregular Margins

One expects the worst performance from the random walk when the row and column sums (margins) are very irregular. This gives rise to sets Σ_{rc} whose convex polyhedral description is not very well-rounded. We show some results from experiments with such "lumpy" tables. So that we can compare with exact results, we use small values of N for which exhaustive enumeration is also feasible.

				r_i
	12	0	0	3
	8	2	2	0
	9	2	1	2
	4	1	2	1
c_j	33	5	5	6
				$N=49$

Figure 5.12: A table with irregular margins.

The table of Figure 5.12 is a table with irregular margins. It was constructed artificially, and is not from an actual study. An exact enumeration took 60.3 user CPU seconds and 1 minute and 2 seconds of real time. There are 252253 tables with the same margins. The χ^2 value for the given table is 9.14, and its exact significance to five significant figures is 0.05388.

The Diaconis-Efron approximation estimates the significance between 0.02696 (without protrusion correction) and 0.04405 (with protrusion correction).

We discuss two experiments on this table. In the first, the walk was run for $T = 2954$ steps to randomize and five runs of $2\tau = 198$ steps each were combined to obtain a median histogram and median significance estimate. This took only .847 user CPU seconds and 2 seconds of real time and gave a significance estimate of .0455. This is within .0016 of the true significance. However, the Kolmogorov-Smirnov distance between the two histograms was much larger, 0.17. This means that significance estimates for some other tables in Σ_{rc} would have been worse. However, the result is well within our bounds for this amount of sampling, which say only that the significance estimate should have been within .44 of the true significance with high probability.

In the second experiment, we used the same number of steps, T , to randomize, but now made five runs of $2000\tau = 198000$ adjacent samples each. According to our bounds, under the eigenvalue hypothesis, this would guarantee that the error in the significance estimate is at most .014 with probability greater than .96. In actuality, we did much better. Our estimate from this experiment to five significant figures is .05388 agreeing with the true value in all these digits. So the error here is less than .000005. Moreover, the Kolmogorov-Smirnov distance between the sample and exact histograms was less than .002.

The time required for this second experiment was 120 seconds of user CPU time, and 2 minutes

37 seconds of real time, or roughly twice the time required for the exact enumeration. The exact and sample histograms are shown in Figure 5.13. It should be noted that the number of samples being computed in each of the five runs was a significant fraction of the number of tables. This makes the accuracy obtained less surprising, however the following is to be noted, here and in the other experiments. If the walk had a convergence rate much worse than our conjectured rate, we might expect to "get stuck" in some small region of the set Σ_{rc} . Instead, our sampling seems to get a representative portion of Σ_{rc} , as we would expect from near-uniform sampling.

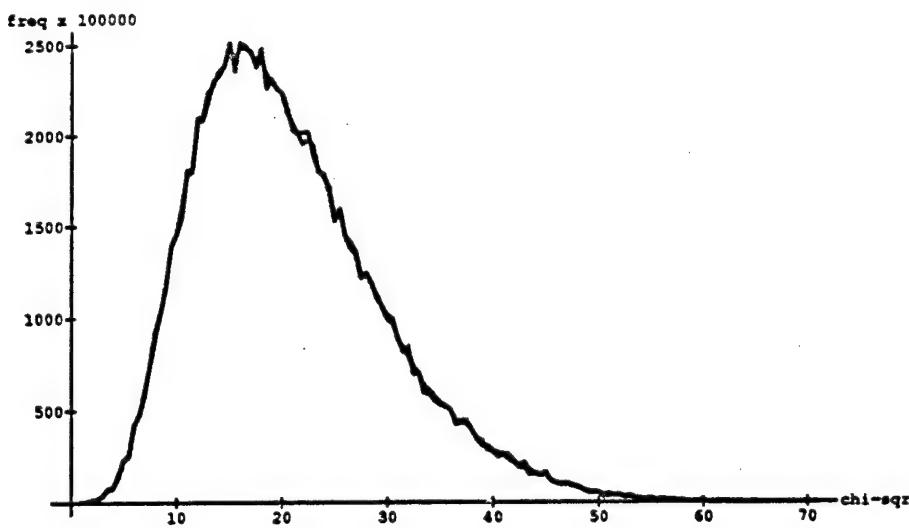
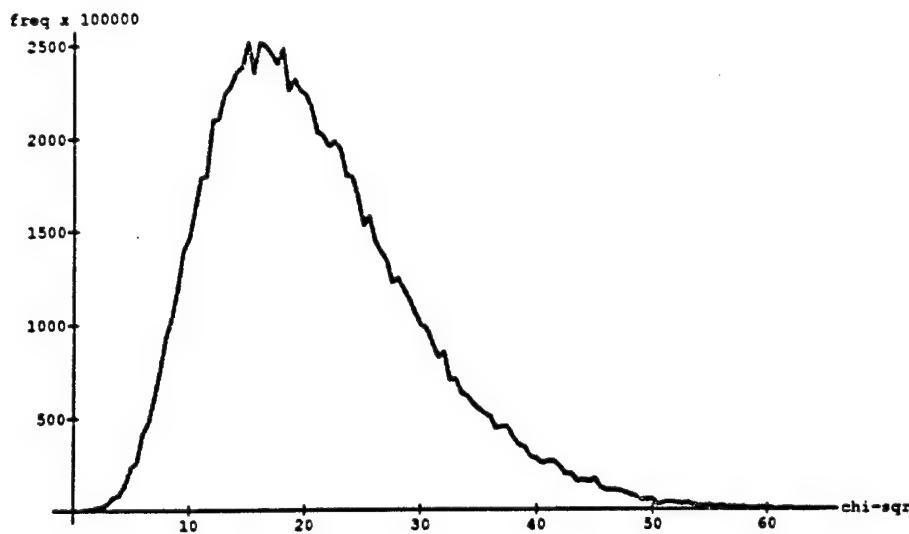


Figure 5.13: Histograms in the second (larger-sample) experiment on the table in Figure 5.12. The sample histogram is depicted alone in the upper plot, and underlying the exact histogram (dark) in the lower plot. Each histogram is normalized to have total mass 100000. Kolmogorov-Smirnov distance between the two histograms is less than 0.002. The significance estimate for the observed table is off by less than 0.000005.

				r_i
	5	2	3	10
	50	7	5	62
	3	6	4	13
	5	3	3	11
	2	7	30	39
c_i	65	25	45	N=135

Figure 5.14: Another table with irregular margins

Figure 5.14 shows another table with irregular margins. The exhaustive enumeration of the corresponding set Σ_{rc} took 54665 CPU seconds, and 15 hours and 21 minutes of real time. There were 239382173 tables in the set. The specified table has $\chi^2 = 72.18$ and exact significance to five significant figures is .76086.

The Diaconis-Efron approximation very accurately estimates the number of tables in Σ_{rc} as 2.33×10^8 , but their approximation grossly overestimates the significance region, resulting in a significance estimate greater than 1, even with their protrusion correction.

We ran our Monte Carlo method, again employing median of sample mean estimates over five runs. In each run we used $T = 49936$ steps to randomize, and took $2000\tau = 2297000$ adjacent samples. The entire process took 23 minutes and 31 seconds, or about 1/30th the time of the exhaustive method.

Our guarantees place our estimate within .027 of the true value with probability greater .96. The resulting significance estimate was .7638, which is off by .003, much less. The histograms are shown in Figure 5.15. The Kolmogorov-Smirnov distance between them is less than .005. In this case, in each of our five runs, the number of samples was only about 1% of the size of the sample space. This result is therefore more convincing than that of the preceding experiment.

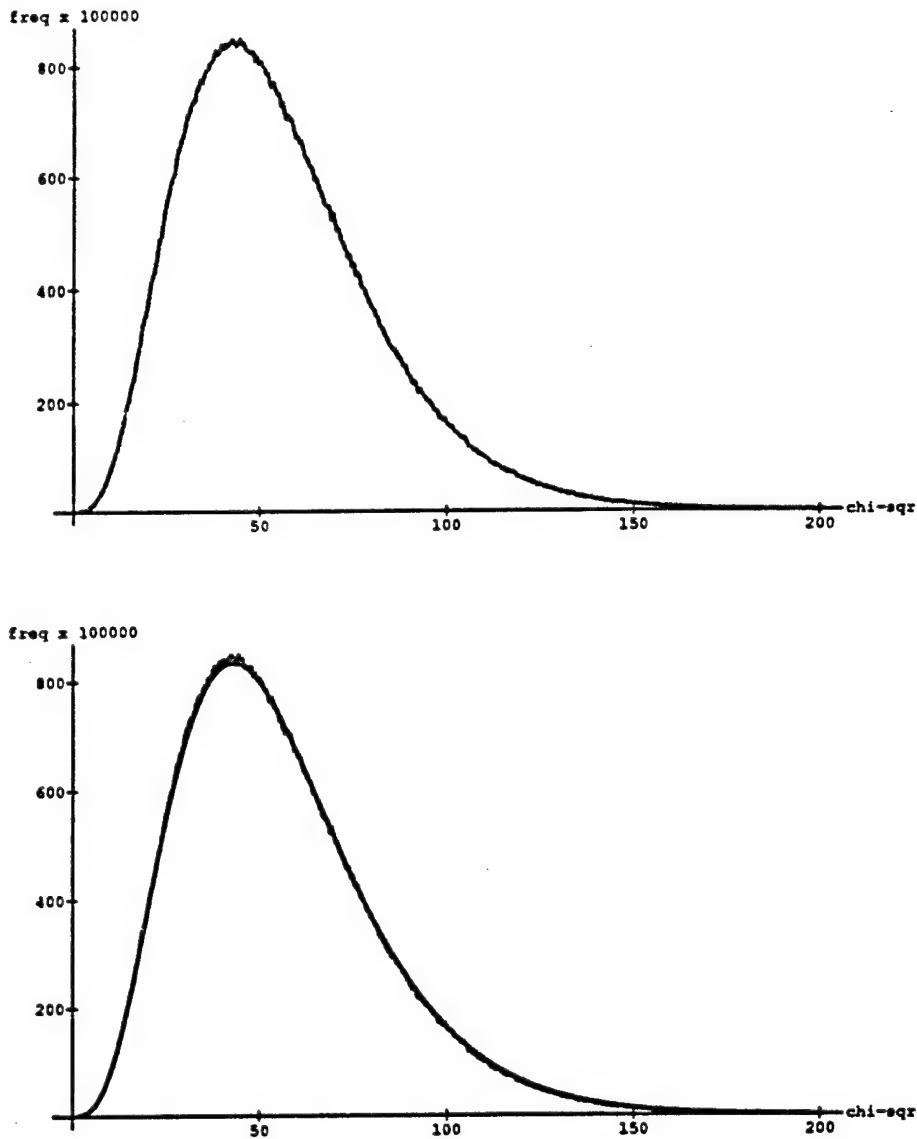


Figure 5.15: Histograms for the table in Figure 5.14. The sample histogram is depicted alone in the upper plot, and underlying the exact histogram (dark) in the lower plot. Each histogram is normalised to have total mass 100000. The Kolmogorov-Smirnov distance between the two histograms is less than 0.005, and the significance estimate for the observed table is off by only 0.003. The exact histogram took over 15 hours to compute, while the sample histogram took under 24 minutes.

5.4.5 Scaling

In the examples above, the large number of steps required to reach stationarity is due essentially to the presence of large margins, which in turn are due to large sample sizes N . One might hope to get good approximations as follows. Scale the table to have grand sum $N' < N$, by multiplying by N'/N throughout, rounding as needed. (One hopes that the effect of rounding is small). Calculate or estimate the significance, for several such N' and extrapolate to the larger N .

The problem with this method is that it is not always easy to extrapolate accurately. Certainly, for very large N , the significance value is not very dependent on N . But gauging the passage from small N' to much larger N , seems to be difficult. Diaconis and Efron attempt such an extrapolation for the Hair-Eye color table of Figure 5.10. Using their Monte-Carlo method, they estimate the significance of the table for $N' = 40$, $N' = 60$, and $N' = 80$, at .036, .051, and .069, respectively. They extrapolate, (they do not say how), to .09 for $N = 592$. We think this is too small.

Eye Color	Hair Color				r_i
	Black	Brunette	Red	Blond	
Brown	23	40	9	2	74
Blue	7	28	6	32	73
Hazel	5	18	5	3	31
Green	2	10	5	5	22
c_j	37	96	25	42	$N'=200$

Figure 5.16: Re-scaled Hair/Eye Color Table.

We experimented with a table in between. For $N' = 200$, the scaled table appears in Figure 5.16. The Diaconis-Efron approximation for the number of such tables is on the order of 10^{11} , so that an exhaustive significance calculation is not feasible. Their analytic approximation with maximum protrusion correction estimates the significance at 0.13. Using our method, we estimate the significance at 0.136. We took $2000\tau = 4816000$ samples in each of five runs, with 133914 steps used to get near stationarity in each run. Assuming the eigenvalue bound holds, and that the actual significance is 0.13, this would have given us accuracy to within .02 with high probability. We therefore believe the exact significance is in the range $0.13 \pm .02$. The experiment took 2975 CPU seconds and just under 53 minutes of real time.

5.5 Provably Polynomial-Time Methods

Under the hypothesized eigenvalue bound, the random walk method we suggest gives a significance estimate within ϵ of the true value with probability at least $1-\beta$, in time polynomial in the dimensions

of the table m and n , the grand sum N , $1/\epsilon$, and $\lg(1/\beta)$. However, the eigenvalue bound is only conjectural.

Using results of Dyer, Frieze, and Kannan [DFK89], one can prove that a modified version of our algorithm converges in polynomial time for a more substantial and interesting subclass of sets Σ_{rc} . Their methods can also be applied to get other provably polynomial-time randomized algorithms for estimating significance under the uniform model, (as well as some other problems that can be solved by sampling, such as approximate counting). In both cases the resulting polynomial bounds, do not give very practical running times. Not only are the degrees of the polynomials substantial, but the constant factors are also large. In practice, one must, as we did, conjecture faster performance to match the empirical performance of our suggested method. So the results in this section do not seem to lead to a better practical solution to the problem.

We will not give a detailed description of their arguments, but we will show how to apply their ideas. For details, the reader will need to refer to their paper. We use some terminology we introduced in Section 5.3, and we need to introduce some more.

For any point $\mathbf{z} = (z_1, z_2, \dots, z_d) \in \mathbb{R}^d$, let $\text{cube}(\mathbf{z})$ denote the d -dimensional unit cube centered at \mathbf{z} ,

$$\text{cube}(\mathbf{z}) = \{y \in \mathbb{R}^d \mid \sum_i |y_i - z_i| \leq 1\}.$$

The integer lattice in d -dimensions is the infinite graph obtained by joining every two points with integer coordinates whose distance is exactly 1. This graph is $2d$ -regular.

Let K be any convex set in \mathbb{R}^d . Consider, now, the finite induced subgraph of the integer lattice whose vertices V are those points v with integer coordinates such that $\text{cube}(v)$ intersects K ,

$$V = \{v \in \mathbb{Z}^d \mid \text{cube}(v) \cap K \neq \emptyset\}.$$

To any vertices $v \in V$, whose degree $\deg v$ is less than $2d$, add $2d - \deg v$ self-loop edges. Call the resulting graph $\text{Latt}(K)$, the lattice graph of K . It is $2d$ -regular, and is connected (by convexity of K). Note that the vertex set V contains two types of points. It contains points in \mathbb{Z}^d that are in K , which we call **interior vertices**, and it also contains points that are not in K , but whose cubes intersect K , which we call **boundary vertices**. Note that if boundary vertices were excluded, it could destroy connectivity.

Call a set in \mathbb{R}^d large and well-rounded if it contains a ball B_I of radius ρ_I and is contained in a ball B_O of radius ρ_O such that $\rho_I \geq 20d^{5/2}$ and $\rho_O/\rho_I \leq \sqrt{d}(d+1)$.

Theorem 5.17 (from [DFK89]) *Let K be a large and well-rounded convex set. Let P be the strongly-aperiodic random walk on $\text{Latt}(K)$. Then P is time-reversible, ergodic, and has uniform stationary distribution on the vertices V , and the second absolute eigenvalue of P satisfies*

$$\lambda_2(P) \leq 1 - 1/4000000d^9\rho_I^4.$$

Proof: That P is time-reversible, ergodic, and symmetric follow from the facts that P is a random walk on a connected regular graph. Scale coordinates in all dimensions by $1/\rho_I$. This puts one in the framework used by [DFK89]: B_I becomes the unit ball, and their value $\delta = 1/\rho_I \leq 1/20d^{5/2}$. Note that this is what they call the *natural* Markov chain, and not the *technical* Markov chain. Now apply their Lemma 1. It gives a bound on the edge-expansion of the graph, which we multiply by $1/4d$ (the minimum probability on any transition of the strongly aperiodic chain) to get a conductance bound. Then square and replace δ with $1/\rho_I$ to obtain this version. \blacksquare

Lemma 5.18 (from [DFK89]) *If K is a large and well-rounded convex set in \mathbb{R}^d , then the ratio of the number of boundary vertices of $\text{Latt}(K)$ to the number of interior vertices of $\text{Latt}(K)$ is at most $3/20d$.*

Proof: This is their Proposition 3, with their value $\eta = \delta \leq 1/20d^{5/2}$. \blacksquare

Now, for a given instance of Σ_{rc} , let $d = (m-1)(n-1)$ and let D be the convex d -dimensional polytope consisting of the set of all $(m-1) \times (n-1)$ tables $X = (X_{ij})$ with nonnegative real-valued entries satisfying

$$\sum_{1 \leq j < n} X_{ij} \leq r_i \text{ for each } i, 1 \leq i < m \quad \sum_{1 \leq i < m} X_{ij} \leq c_j \text{ for each } j, 1 \leq j < n \quad (5.3)$$

The interior vertices of $\text{Latt}(D)$ correspond precisely with tables in Σ_{rc} . The random walk on the graph $\text{Latt}(D)$ is very similar to our basic random walk (Algorithm 5.2). To simulate the walk on $\text{Latt}(D)$, we add 1 or subtract 1 from some entry of the table T_{ij} with $1 \leq i < m$ and $1 \leq j < n$. This corresponds to a move of Algorithm 5.2 in which $i_1 = i$, $j_1 = j$, $i_2 = m$, $j_2 = n$, except that now we allow the table T_{ij} to be ‘slightly outside’ of Σ_{rc} . (Changing all entries by at most 1, one will be able to get to a table within Σ_{rc} .)

Call an instance of Σ_{rc} large and well-rounded if the associated d -dimensional polytope D is large and well-rounded. The following corollary is almost immediate.

Theorem 5.19 (Near-Uniform Generation in Large Well-Rounded Σ_{rc}) *There is an algorithm which, given any large and well-rounded instance of Σ_{rc} , an ϵ and δ (with $0 < \epsilon < 1$ and $0 < \delta < 1$), runs in time polynomial in m , n , N , and $\lg 1/\epsilon$, $\lg 1/\delta$. At its termination, with probability at least $1 - \delta$, the algorithm reports a ‘success’ and outputs a table in Σ_{rc} according to a distribution that is within ratio ϵ of the uniform distribution. Otherwise it reports a ‘failure’ and outputs no table.*

Proof: Use the random walk to draw samples from the vertices of $\text{Latt}(D)$, within ratio ϵ of uniform. It is clear that we can simulate steps of the walk time polynomial in m , n , and N . Theorem 5.17 and Theorem 1.8, together the facts that we must have $\rho_I \leq N$, and $|\Sigma_{rc}| \leq N^d$,

imply that $O(N^4d^9(d \ln N + \lg(1/\epsilon)))$ steps suffice to get a vertex distributed within ratio ϵ of uniform on $\text{Latt}(D)$. We reject the sample if it is a boundary vertex. By Lemma 5.18 and the near-uniformity of the sample, this happens with probability at most $\frac{3}{20d}(1+\epsilon) < 3/10$. The result follows by repeating $O(\lg 1/\delta)$ times, stopping and reporting a 'success' if a sample is an interior vertex, and outputting the corresponding table in Σ_{rc} . We report 'failure' if none of the $O(\lg 1/\delta)$ samples are interior vertices. \square

It follows that we can use this alternative technique in our suggested method to get good estimates for significance or any other mean-value quantity.

Corollary 5.20 *For large well-rounded instances of Σ_{rc} there is an algorithm that given any table $T \in \Sigma_{rc}$ runs in time polynomial in within m , n , and N , $1/\epsilon$ and $\lg 1/\delta$ and outputs an estimate M for the volume-based significance p of $\chi^2(T)$ such that $|M - p| \leq \epsilon$ with probability at least $1/\delta$.*

Large well-rounded instances of Σ_{rc} correspond to instances in which the grand sum is large and the margins are quite regular. The following example illustrates the most regular case. Other classes of well-rounded instances can be established in a similar way.

Example 5.21 (Magic Squares) The set H_{nr} , magic squares of order n and margins r , is the set of square arrays in which every row sum and column sum is r . (Some authors reserve the term 'magic squares' for arrays in which the main diagonals also sum to r .) This is a well-studied class of instances of Σ_{rc} . (See [Sta86, p. 232] and [GC77].) It is not hard to see that for $r \geq 40(n-1)^7$ these instances are all well-rounded. Let $d = (n-1)^2$. The d -dimensional polytope D for H_{nr} contains the d -dimensional cube $[\frac{(n-2)r}{d}, \frac{r(n-1)}{d}]^d$. It follows that it contains a ball of radius at least half the width of the cube, $\rho_I = \frac{r}{2d}$. Moreover, D is clearly contained in a ball of radius $\rho_O = rn^{-1/2}$. The ratio of these two is $\rho_O/\rho_I \leq 2d^{3/4}$, and so it follows that for $r \geq 40d^{7/2} = 40(n-1)^7$ that D is large and well-rounded. From the algorithm above, we will be able to get nearly uniform random magic squares in roughly $O(r^4d^6 \ln r)$ steps of the walk. Once near uniformity, one can obtain 'good' samples in $O(r^4d^5)$ steps using the ideas in Chapter 3. \square

Remark: More generally, the algorithm described in [DFK89] can be used to estimate significance in cases that are not well-rounded. The results of [DFK89] can be used to approximate the number of tables in Σ_{rc} and the number of tables in a given significance region as follows: D is convex but not well-rounded, there is an affine transformation of \mathbb{R}^d that takes D to a large well-rounded convex set K , and this transformation is polynomial-time computable [DFK89, GLS88]. However, the image in K of the integer lattice points in D will form a grid in K having unequal (but regular) spacings in the various coordinates. Still, one can approximate the number of such lattice points within K (or a given region of K) by using a finer lattice in which the spacings are equal in all coordinates. This in turn can be used to approximate the number of lattice points in D . \square

5.6 More on the Eigenvalue Hypothesis

Our eigenvalue hypothesis was based on Payne and Weinberger's Theorem (5.7) for the Laplacian in the continuous setting. Because of the similarity of the Rayleigh quotients involved in both cases, it is natural to believe that the eigenvalues of the Laplacian of the random walk on the lattice within a (large enough) body K will be approximately that of the body. (Increasing the scale of the body is equivalent to placing a finer mesh within the original body.)

I believe that the following statement or one of essentially similar form holds generally for lattice graphs within large (not necessarily well-rounded) convex sets.

Conjecture 5.22 (Lattice Approximation Hypothesis) *There exists a (low-degree) polynomial $r(d)$ and a constant c such that if K is any body in \mathbb{R}^d containing a ball of radius $\rho_1 > r(d)$, the following holds for $\text{Latt}(K) = (V, E)$.*

If $\phi : V \rightarrow \mathbb{R}$ is any function with $\sum_{v \in V} \phi^2(v) = 1$ and $\sum_{v \in V} \phi(v) = 0$, then there is a function f on K with bounded second-derivative such that

$$\int_K f^2 dK = 1 \quad \int_K f dK = 0$$

and

$$\int_K |\nabla f|^2 dK \leq c \sum_{e=\{x,y\} \in E} (\phi(y) - \phi(x))^2.$$

If this were true, one could prove the following useful result.

Theorem 5.23 *Suppose the preceding conjecture holds, and let K be a convex body in \mathbb{R}^d with diameter Δ . Let $G = (V, E) = \text{Latt}(K)$ let P be the natural random walk on G , and let $L = I - P$ be the associated Laplacian. Then*

$$\mu_1(L) \geq \frac{\pi^2}{2dc\Delta^2}.$$

Proof: Let ϕ be an eigenfunction of $L = I - P$ corresponding to $\mu_1(L)$, normalized it so that $\sum_v \phi^2(v) = 1$. This satisfies the preconditions of the conjecture above, so let f be a corresponding function on K guaranteed by the conjecture. This function is admissible in the Payne-Weinberger lower bound (5.1), which gives us

$$\sum_{e=\{x,y\} \in E} (\phi(y) - \phi(x))^2 \geq \frac{1}{c} \int_K |\nabla f|^2 dK \geq \frac{\pi^2}{c\Delta^2}.$$

Therefore, by (1.8), we have

$$\mu_1 = \frac{1}{2d} \sum_{e=\{x,y\} \in E} (\phi(y) - \phi(x))^2 \geq \frac{\pi^2}{2dc\Delta^2}.$$

■

The theorem would imply $O(n\Delta^2 \ln |V|) \subseteq O(n^2\Delta^2 \ln \Delta)$ convergence for walks on connected lattice graphs within sufficiently large convex bodies. The conjecture could be weakened slightly to allow c to depend as a small polynomial on the dimension n , with a corresponding weaker theorem.

Remark: Dyer, Frieze, and Kannan use an isoperimetric inequality combined with Jerrum and Sinclair's conductance bounds (see Theorem 1.16) and it seems that the resulting squaring gives them an eigenvalue bound that depends on Δ essentially as $1/\Delta^4$ for fixed dimension. We know that if K is a d -dimensional rectangular box the eigenvalue depends on Δ as $1/\Delta^2$.

Lovász and Simonovits [LS90] give an isoperimetric inequality for similar lattice graphs, that almost proves a bound of this form. Unfortunately, they are forced to neglect sets whose size falls below a certain value $m(D)$. This arises because of problems with the isoperimetric inequality at sets of points taken near the boundary. They bound the size of the boundary by $m(D)$ and neglect them in their inequality. The resulting isoperimetric inequality does not give them provably rapid mixing to zero variation distance, but to a distance that depends on the maximum total variation of the initial distribution π_0 from the stationary distribution π on sets of sets of size $m(D)$. They show that they can still use this bound in their application.

We are suggesting that it may be better to handle the problem directly as a variational problem, instead of using an isoperimetric inequality. It may be the case that only a weaker isoperimetric inequality (as that of Lovász and Simonovits) holds for these graphs. \square

Chapter 6

Directions for Future Work

There are several areas that seem promising to investigate in the near future. Here are some ideas to consider that relate to the work in this thesis.

1. Many of the beautiful convergence results are based on isoperimetric inequalities that had arisen earlier in the study of the Laplacian on Riemannian manifolds. Much of the framework used in Chapter 1 was based on geometric intuitions from the continuous setting. Almost certainly, there are more connections to explore. A solid background in analysis is probably required. I felt that this was a problem in resolving the approximation conjecture in the last chapter. That result, as well as others relating these two branches of mathematics would be extremely interesting.
2. In the last couple of years many authors have recognized the remarkable prospects for using expander graphs to reduce randomness requirements in probabilistic algorithms. In Chapter 4, we proved some new independence-like properties of correlated samples obtained from expanders, discussed some properties that others had proved, and showed a new way in which they could be combined. What further properties do such samples nearly share with truly independent samples? Good results in this arena would have implications in many probabilistic algorithms.
3. Most of the tight convergence bounds for Markov chains come from full spectrum bounds, which in turn come from harmonic analysis. When can walks on graphs be lifted to walks on groups? This question was raised and partially answered by Diaconis and Shahshahani [DS87]. We discussed some such cases in Chapter 2, with the purpose of showing some general diameter-based bounds for such walks, even when the harmonic analysis does not seem tractable. These are nice in the absence of other bounds, but when one can do the harmonic analysis one does better. Often one can't. So is there anything in-between? For example, can one sometimes

use symmetry and knowledge of a particular group to construct a system of paths that gives a convergence bound that is $o(\Delta^2 \ln |V|)$? What about coupling methods? These questions seem hard. The advantage to the interested investigator is that there are some nice examples that have been treated with tight results, and the background material is well-developed. See [Dia88].

4. A matroid in S is a family \mathcal{I} of subsets $I \subset S$ that is closed under containment (\mathcal{I} an ideal), and such that the cardinality of every maximal subset I is the same. The Bernoulli-Laplace process that we studied in Chapter 2 can be viewed as a random walk on the basis-exchange graph of the trivial matroid whose bases (maximal independent sets) are all k -sets of $[n]$. Broder's random walk on matchings [Bro86, JS88] may be viewed as a walk on the top two layers of the graph of the associated matroid. Are there any general conductance bounds for such walks? Intuitively, ideals in the hypercube are 'convex' and should not have conductance much worse than the cube. Mihail and Vazirani [MV88] have posed conjectures about the conductance of similar graphs. A geometric technique seems promising. This question for the spanning-tree matroid has still not been answered despite significant efforts by several people. Answers could be applied to sampling in matroids.
5. To what other problems can the existing eigenvalue bounds be applied? It seems like one reasonable place to look is in the field of numerical analysis, and at iterative matrix methods in particular.

Appendix A

Notions of Approximation

A.1 Point Approximations

For $\hat{a}, a > 0$ and $r \geq 1$, we say that \hat{a} approximates a to within ratio r if

$$\hat{a}/r \leq a \leq \hat{a}r.$$

Note that we require all of the quantities involved to be *positive*. This form of measuring error is essentially equivalent to the usual notion of relative error but more easily handled in combinatorial settings. Recall that one says \hat{a} approximates a to within relative error $\epsilon < 1$ if

$$a(1 - \epsilon) \leq \hat{a} \leq a(1 + \epsilon).$$

If \hat{a} approximates a within ratio $(1 + \epsilon)$, with $\epsilon < 1$, then \hat{a} approximates a to within relative error ϵ . Almost conversely, if \hat{a} approximates a within relative error $\epsilon < 1$, then \hat{a} approximates a within ratio $1/(1 - \epsilon) = 1 + \epsilon + \epsilon^2 + \dots$, which is at most $1 + 2\epsilon$ if $\epsilon \leq 1/2$.

Here are some basic properties of approximation within ratio that the reader should keep in mind. Suppose that \hat{a} approximates a and \hat{b} approximates b , each within ratio r . Then $\hat{a} + \hat{b}$ also approximates $a + b$ within ratio r . We call this the *sum rule* for approximations within ratio. By induction, it holds for all finite sum of positive summands. There is the following similar product rule. If \hat{a} approximates a within ratio r and \hat{b} approximates b within ratio r' then $\hat{a}\hat{b}$ approximates the product ab within ratio rr' . A useful corollary is that if $\hat{a}_1, \hat{a}_2, \dots, \hat{a}_n$ approximate a_1, a_2, \dots, a_n , respectively, each within ratio $1 + \frac{\epsilon}{2^n}$ for $\epsilon < 1$, then the product $\prod_i \hat{a}_i$ approximates the product $\prod_i a_i$ within ratio $(1 + \frac{\epsilon}{2^n})^n \leq (1 + \epsilon)$.

A.2 Approximate Distributions

To be able to say with precision that two distributions are “close” to each other, we utilize various measures of distance between distributions: (total) variation, and separation. Here we state their definitions and some of their properties.

A.2.1 Total Variation

Variation distance or total variation is a common statistical measure of distance between two distributions based on the \mathcal{L}^1 norm.

If P and Q are two probability distributions on a finite set S , the variation distance or total variation between P and Q , denoted $\|P - Q\|$, is defined by:

$$\|P - Q\| = \frac{1}{2} \sum_{s \in S} |P(s) - Q(s)|.$$

Some authors define variation without the $1/2$ factor.

Proposition A.1 *Variation distance satisfies the following properties:*

1. $0 \leq \|P - Q\| \leq 1$;
2. $\|P - Q\|$ is a metric: it is nonnegative; it is zero if and only if $P = Q$; and it satisfies the triangle inequality;
3. $\|P - Q\| = \max_{A \subseteq S} |P(A) - Q(A)|$;
4. $\|P - Q\| = P(A) - Q(A)$, where $A = \{s \in S \mid P(s) > Q(s)\}$;
5. $\|P - Q\| = \frac{1}{2} \sup_{f: \|f\|_\infty \leq 1} |\mathbb{E}_P[f] - \mathbb{E}_Q[f]|$, where f is a real-valued function on S .

Proof: (Outline) (1) Easy. (2) This follows from metric properties of the \mathcal{L}^1 distance. (3) (4) (5) These follow by separating S into the three classes: $A = \{s \mid P(s) > Q(s)\}$, $B = \{s \mid P(s) = Q(s)\}$, and $C = \{s \mid P(s) < Q(s)\}$. Show that A simultaneously gives the maximum and equality with the definition, yielding (3) and (4). To get (5) define a function f that is 1 on elements in A , takes any values on elements in B , and -1 on elements in C , and show that this yields the stated supremum and equality with the definition. ■

A.2.2 Separation and Relative Pointwise Distance

Relative pointwise distance is also a measure of discrepancy between distributions, but is not a metric. We define the relative pointwise distance from P to Q as

$$\text{rpd}(P, Q) = \max_{s \in S} \frac{|P(s) - Q(s)|}{Q(s)},$$

in other words, the maximum relative error at any point. This distance is not symmetric. Generally, Q is taken to be some fixed distribution, and we talk about the distance of varying distributions P from the fixed Q . Notice that this distance is only defined when $Q(s) > 0$ for all $s \in S$. In this work, Q is generally the stationary distribution of an ergodic Markov chain, and such a distribution is necessarily positive everywhere on the state space.

Closely related to relative pointwise distance is the separation of P from Q , denoted $\text{sep}(P, Q)$. It is defined by:

$$\text{sep}(P, Q) = \max_{s \in S} \frac{Q(s) - P(s)}{Q(s)}.$$

Note the absence of the absolute value. Separation arises naturally in the study of strong stationary times. [AD86]

Proposition A.2 *Separation and relative pointwise distance satisfy the following properties:*

1. $\text{rpd}(P, Q) \geq \text{sep}(P, Q) \geq 0$;
2. $\text{rpd}(P, Q) = 0$ (also $\text{sep}(P, Q) = 0$) if and only if $P = Q$;
3. $\text{sep}(P, Q) \leq 1$, with equality if and only if $P(s) = 0$ for some $s \in S$;
4. $\|P - Q\| \leq \frac{1}{2} \text{rpd}(P, Q)$
5. $\|P - Q\| \leq \text{sep}(P, Q)$
6. if U is the uniform distribution, then $\text{rpd}(P, U) \leq |S| - 1$.

Proof: (Outline) (1)(2)(3) Easy. (4) By definition, for all $s \in S$ we have $|P(s) - Q(s)| \leq \text{rpd}(P, Q)Q(s)$. Sum both sides of this inequality over $s \in S$ to get: $\sum_s |P(s) - Q(s)| \leq \text{rpd}(P, Q) \sum_s Q(s)$. Then divide by two to get the desired result. (5) Let C be the set defined in the proof of item 4 of Proposition A.1. Then we have

$$\|P - Q\| = Q(C) - P(C) = \sum_{s \in C} Q(s) \frac{Q(s) - P(s)}{Q(s)} \leq \text{sep}(P, Q) \sum_{s \in C} Q(s) \leq \text{sep}(P, Q).$$

(6) Again, easy. ■

A.2.3 Approximation within Ratio

For two distributions P and Q supported everywhere on S , we say that P approximates Q within ratio r if for all points $s \in S$, $P(s)$ approximates $Q(s)$ within ratio r pointwise.

Our earlier comments relating relative error and approximation within ratio for single points carry over here. In particular, if for some positive ϵ , P approximates Q within ratio $(1 + \epsilon)$, then $\text{rpd}(P, Q) \leq \epsilon$. Almost conversely, if $\text{rpd}(P, Q) \leq \epsilon/2$, and $\epsilon \leq 1$, then P approximates Q within ratio $1 + \epsilon$.

The following two properties are fundamental.

Proposition A.3 *If distribution P approximates Q within ratio r on a finite set V , then for every subset $A \subseteq V$, $P(A) = \sum_{v \in A} P(v)$ is within ratio r of $Q(A)$.*

Proof: Immediate from the sum rule. \square

Proposition A.4 *Let X and Y be drawn from a finite set V with distributions P and Q , respectively. Suppose that P approximates Q within ratio r . Then for any function b on V , the distribution of $b(X)$ also approximates the distribution of $b(Y)$ within ratio r .*

Proof: For any value z in the range of b , consider the probability $\Pr\{b(X) = z\}$. We wish to show that this is within ratio r of the probability $\Pr\{b(Y) = z\}$. For this, simply consider the set V_z of $v \in V$ such that $b(v) = z$. Now, $\Pr\{b(X) = z\} = \Pr\{X \in V_z\} = P(V_z)$. Applying the preceding proposition, we find that this is within ratio r of $Q(V_z) = \Pr\{Y \in V_z\} = \Pr\{b(Y) = z\}$. Since this holds for each z in the range of b , we have proved the theorem. \square

A.2.4 Kolmogorov-Smirnov Distance

In reporting our experiments in Chapter 5 we use a distance measure that we call Kolmogorov-Smirnov distance. We explain that notion here.

Recall that if X is a real-valued random-variable, the function $F(z) = \Pr\{X \leq z\}$ is called its cumulative distribution function. If $F(z)$ and $G(z)$ are the cumulative distribution functions for two random variables, the Kolmogorov-Smirnov distance (K-S distance) between them is $\sup_{z \in \mathbb{R}} |F(z) - G(z)|$.

Since the significance level of a given 'observed' value z_0 under the 'null hypothesis' that it was drawn according to the distribution F is $F(z_0)$, the K-S distance gives the maximum absolute difference between any two significance levels for the same observed value, one measured under F and the other under G . So, for example, if F is an approximation we construct for G , and the main purpose of the approximation is to obtain approximations to significance levels under G , the K-S distance is the natural distance measure to use, and this is why it is applied in Chapter 5.

The quantity we report as K-S distance is actually a discretized version of K-S distance. If f is a histogram, let $F(b)$ denote the cumulative sum over all bins up to and including bin b . If F and G are the cumulative sums for two histograms with common bin sizes, the quantity we report as K-S distance is the maximum value of $|F(b) - G(b)|$ over all bins b .

Appendix B

Sampling from Near-Stationary Chains

In Chapter 1 we showed a number of techniques for obtaining convergence guarantees. In this appendix we prove some results showing how to use convergence guarantees to yield guarantees about the quality of statistics based on samples from the nearly stationary Markov chain.

In Section B.1, we show that well-spaced nearly stationary samples from a Markov chain yield nearly independent, nearly stationary samples and that statistics based on these samples will nearly match, in distribution, the statistics based on independent and identically distributed (i.i.d.) samples drawn according to the stationary distribution.

Recall that in Chapter 3 we discussed mean value estimation using highly correlated closely spaced samples drawn from a Markov chain. There we assumed that the chain started *exactly* in the stationary distribution. This was fine for the applications described in Examples 3.7 and 3.8 and in Chapter 4, where it was easy to get an initial sample drawn according to the stationary distribution. However, often one may not have an efficient method of starting the chain exactly in the stationary distribution. In fact, the Markov chain may be the only known means of generating even nearly-uniform samples efficiently. (This is the case, for example, in Example 3.9 and in Chapter 5). In these cases one wants first to simulate the chain until one knows, by a convergence guarantee, that the chain is *nearly* stationary, and then one wants to begin sampling. In Section B.2, we show how to extend the error bounds in Corollary 3.5 to this situation of sampling from the near-stationary chain. We also explain how the Median Lemma (Lemma 3.1) can be applied with such estimates, provided that between successive estimates we take enough steps to insure their near-independence.

We work here with the notion of 'approximation within ratio.' Two important combinatorial properties of approximation within ratio, the sum and product rule, are discussed in Appendix A. The reader should be familiar with that material before proceeding.

B.1 Nearly Independent, Near-Stationary Samples

When samples are drawn from a Markov chain, we may not be able to claim that each successive sample is independent or even exactly uniformly distributed. However, we will typically have guarantees that regardless of the values of prior samples, each successive sample has, to within a small tolerance, approximately the stationary distribution. In this case, we can guarantee that they share important approximation properties with true i.i.d. samples.

Suppose that Y_1, Y_2, \dots, Y_n are each drawn randomly from V . We will say that they are jointly iid Q within ratio $1 + \epsilon$ if the joint distribution of Y_1, Y_2, \dots, Y_n approximates the product distribution Q^n of n independent Q -distributed random variables within ratio $1 + \epsilon$. The following theorem tells us that well-spaced samples drawn from a Markov chain are jointly i.i.d. π within ratio $1 + \epsilon$.

Theorem B.1 *Let $\{X_k \mid k \geq 0\}$ be a time-reversible ergodic Markov chain with a convergence guarantee $T(\epsilon)$.*

If $t > T(\epsilon/2n)$ and $\{X_{ti} \mid 1 \leq i \leq n\}$ are n samples drawn every t steps from the chain. Then these samples are jointly i.i.d. π within ratio $1 + \epsilon$.

Proof: Let

$$P_i(x; z_1, z_2, \dots, z_{i-1}) = \Pr\{X_{ti} = x \mid X_t = z_1, X_{2t} = z_2, \dots, X_{t(i-1)} = z_{i-1}\}.$$

Then we have

$$\Pr\{X_t = z_1, X_{2t} = z_2, \dots, X_{nt} = z_n\} = \prod_{1 \leq i \leq n} P_i(x; z_1, z_2, \dots, z_{i-1}).$$

Now, by the Markov property, we have

$$P_i(x; z_1, z_2, \dots, z_{i-1}) = \Pr\{X_{ti} = x \mid X_{t(i-1)} = z_{i-1}\}$$

and by the convergence guarantee of $T(\epsilon/2n)$, this approximates the distribution π within ratio $1 + \frac{\epsilon}{2n}$. By the product rule for approximations within ratio, the product $\prod_i P_i(x_i; z_1, z_2, \dots, z_{i-1})$ approximates the product of $\prod_i \pi(z_i)$ within ratio $1 + \epsilon$. (See Appendix A.) \blacksquare

The next theorem tells us that statistics based on such nearly i.i.d. samples are close, in distribution, to the corresponding statistics based on true i.i.d. samples.

Theorem B.2 *Let b be any function on V^n . If $\{X_i \mid 1 \leq i \leq n\}$ are jointly i.i.d. π within ratio $1 + \epsilon$, then the distribution of $b(X_1, X_2, \dots, X_n)$ is within ratio $1 + \epsilon$ of the distribution of $b(Y_1, Y_2, \dots, Y_n)$, where Y_1, Y_2, \dots, Y_n are i.i.d. samples from V with common distribution π .*

Proof: Since the X_i are jointly i.i.d. π within ratio $(1 + \epsilon)$, by definition the distribution of (X_1, X_2, \dots, X_n) on V^n is within ratio $(1 + \epsilon)$ of the distribution of (Y_1, Y_2, \dots, Y_n) . Apply Proposition A.4 from Appendix A. \blacksquare

B.2 Other Near-Stationary Samples

In Chapter 3 we introduced the mean value estimator C_n^t , which was the sample mean based on n samples drawn some t steps apart from a Markov chain that was evolving *exactly* in the stationary distribution. Here we show how the error bounds we proved there hold approximately when we draw the samples from a nearly stationary chain. We assume the reader has already studied the material in Chapter 3.

Let P be an ergodic Markov chain with stationary distribution π on the finite state space V . Take any $t \geq 1$, and consider samples $X_0, X_t, X_{2t}, \dots, X_{(n-1)t}$ drawn t steps apart from the chain, where π_0 is the distribution of the first sample X_0 . These n samples may be viewed collectively as a sample drawn from V^n according to the joint probability

$$\Pr\{(X_0, X_t, X_{2t}, \dots, X_{(n-1)t}) = (x_0, x_1, x_2, \dots, x_{n-1})\} = D(\pi_0; x_0, x_1, x_2, \dots, x_{n-1}),$$

where

$$D(\pi_0; x_0, x_1, x_2, \dots, x_{n-1}) \stackrel{\text{def}}{=} \pi_0(x_0)P^{(t)}_{x_0 x_1}P^{(t)}_{x_1 x_2}P^{(t)}_{x_2 x_3} \cdots P^{(t)}_{x_{n-2} x_{n-1}}$$

Now, if the distribution π_0 approximates π to within ratio r , then the value

$$D(\pi_0; x_0, x_1, x_2, \dots, x_{n-1})$$

will approximate

$$D(\pi; x_0, x_1, x_2, \dots, x_{n-1})$$

within ratio r . (Note: $D(\pi; x_0, x_1, x_2, \dots, x_{n-1})$ is the distribution of n samples drawn t steps apart from the chain that is started *exactly* in the stationary distribution π .) By the sum rule, this extends to subsets; if $S \subseteq V^n$ is any subset of states, then to within ratio r , the probability

$$\Pr\{(X_0, X_t, X_{2t}, \dots, X_{(n-1)t}) \in S\} = \sum_{(x_0, x_1, \dots, x_{n-1}) \in S} D(\pi_0; x_0, x_1, x_2, \dots, x_{n-1})$$

that our sequence of samples lies in S approximates the probability

$$\sum_{(x_0, x_1, \dots, x_{n-1}) \in S} D(\pi; x_0, x_1, x_2, \dots, x_{n-1})$$

of the same event had we started the chain in the stationary distribution π .

From this, the next theorem follows immediately.

Theorem B.3 *Let P be an ergodic Markov chain with stationary distribution π on V . Let X_0 be drawn according to π_0 , and let $X_0, X_t, X_{2t}, \dots, X_{(n-1)t}$ be n samples drawn t steps apart from the chain starting at X_0 . Similarly, let Y_0 be drawn according to π , and let $Y_0, Y_t, Y_{2t}, \dots, Y_{(n-1)t}$ be n samples drawn t steps apart from the chain starting at Y_0 . If π_0 approximates π within ratio r , then the distribution of $(X_0, X_1, \dots, X_{n-1})$ on V^n approximates the distribution of $(Y_0, Y_1, \dots, Y_{n-1})$ to within the same ratio r .*

Proposition B.2 can then be used to conclude that that sampling from the near-stationary chain gives statistics close (in distribution) to sampling from the stationary chain. For example, we can extend Corollary 3.5 as follows.

Corollary B.4 (Chebyshev Bounds in the Near-Stationary Case) *Adopt the setup of Theorem 3.3 and its corollaries for the mean-value estimator C_n^t (with t odd) on a time-reversible ergodic Markov chain. However, suppose that X_0 is drawn, not from the stationary distribution π , but instead according to a distribution π_0 that approximates π within ratio $(1 + \epsilon)$ ($\epsilon \geq 0$). Then for $n \geq (2\tau(t) - 1)k$, and t odd, we can guarantee*

$$\Pr\{|C_n^t - h_1| > \beta\} \leq \frac{h_2\beta^2}{k}(1 + \epsilon).$$

In particular if $n \geq (8\tau(t) - 4)h_2(1 + \epsilon)/\beta^2$, then

$$\Pr\{|C_n^t - h_1| > \beta\} \leq \frac{1}{4}.$$

Proof: View C_n^t as a function on the finite set V^n . We apply the same reasoning as in Proposition A.4 together with the preceding theorem. (In fact, combining these immediately gives the theorem, but we explain how the proof of the Proposition works in this particular setting.)

Let E be the event that $|C_n^t - h_1| > \beta$. This event occurs if $(X_0, X_t, X_{2t}, \dots, X_{n-1})$ falls in a certain subset $S_E \subseteq V^n$ of possible sequences of samples. By the previous theorem and Proposition A.3, the probability that $(X_0, X_t, X_{2t}, X_{n-1})$ falls in the subset S_E is within ratio $(1 + \epsilon)$ of the probability of what it would have been were X_0 drawn exactly stationary. We gave an upper bound on the probability of the event E in this stationary case in Corollary 3.5. We get the new bound by multiplying that probability by $(1 + \epsilon)$. This proves the theorem. \blacksquare

Remark: Note that, when sampling from the near-stationary chain, the estimator C_n^t is no longer unbiased. We are not applying a Chebyshev bound to C_n^t in the near-stationary case. Instead, having proved a bound on the probability of the event E in the stationary case, we can apply the preceding Theorem to bound the probability of the same event in the near-stationary case. This approach seems altogether much simpler than trying to reason out a similar bound directly using error bounds for the mean and variance of C_n^t in the near-stationary case. \square

B.3 On Near-Independence and the Median Lemma

In Chapter 3, we discussed the Median Lemma (Lemma 3.1, page 54) as an efficient way to use repetition to reduce the probable error of an estimate. That lemma relies on certain independence properties, and it may not immediately be clear that this lemma can be applied with the estimator C_n^t , which uses correlated samples.

It is important to understand the type of independence required. Notice, in particular, that the proof does not require that the *samples* on which the m estimates α_i are based be independent; we only use the fact the events

$$E_i = \{|\alpha_i - a| > \beta\}$$

occur with probability bounded by $1/4$ independent of the previous events E_i . In other words, we require that for $1 \leq i \leq m$ we have

$$\Pr\{E_i \mid E_1, E_2, \dots, E_{i-1}\} \leq 1/4. \quad (\text{B.1})$$

The law of conditional probability then gives Lemma 3.1.

So the use of correlated samples in the estimator C_n^i does not immediately preclude the application of the Median Lemma with repeated trials of that estimator. We only need to guarantee that this independence condition holds for the separate trials. The following algorithm does this, summarizing the method suggested by the discussion here and in the previous section.

Algorithm B.5 (Estimation from the Near-Stationary Chain) Let $T(\epsilon)$ be a convergence guarantee for the reversible ergodic Markov chain P . Adopting the remainder of the setup of Chapter 3, consider the following procedure. Let ϵ be any nonnegative real number.

Start the chain in any initial state.

for $1 \leq i \leq m$ do begin

 Run the chain P for $T(\epsilon)$ steps.

 Starting with this sample X_0 , compute C_n^i based on $n = \lceil (8\tau(t) - 4)h_2(1+\epsilon)/\beta^2 \rceil$ samples drawn t steps apart, and call the resulting estimate α_i .

Compute the median M of the m values α_i , $1 \leq i \leq m$. Output M .

Theorem B.6 *The estimate M produced by Algorithm B.5 satisfies*

$$\Pr\{|M - h_1| > \beta\} \leq 2^{-m}.$$

Proof: (Outline) Corollary B.4 provides a bound of $1/4$ on the probability of each event $E_i = \{|\alpha_i - a| > \beta\}$. Combining the Markov property (1.1) with the definition of a convergence guarantee, it is straightforward to arrive at the independence condition (B.1) between the events E_i . \blacksquare

Appendix C

Enumerating Contingency Tables

This appendix provides a brief review of some methods for counting Σ_{rc} and some interesting combinatorial interpretations of Σ_{rc} . We discuss various formulas for counting the set using groups, symmetric functions, recurrences, and generating functions. We prove the $\#P$ -completeness of a similar set, suggesting that the problem of counting Σ_{rc} is hard. (The complexity of counting Σ_{rc} , however, remains unknown.) In the last section, we outline how the techniques of Chapter 5 can be extended to apply to approximately counting Σ_{rc} .

Recall from Chapter 5, our usual interpretations of the notations: r , c , N , Σ_{rc} , m , and n . In addition, we use the notation M_{rc} in this appendix to denote the cardinality $|\Sigma_{rc}|$.

C.1 Classical Counting Approaches

C.1.1 Exhaustive Enumeration

Though Σ_{rc} is in general very large, one can sometimes afford to enumerate its elements exhaustively. This, of course, affords a means of exactly calculating the cardinality M_{rc} as well as various other functionals on the set, such as mean values of functions and significance levels under the uniform distribution. A number of authors [Mar72] [BW73] [Han74] have suggested exhaustive algorithms. These work by stepping through the set Σ_{rc} , one table at a time. They begin at some canonically constructed initial table, and proceed by making small changes to the cell entries, so that the tables increase monotonically in some linear ordering. We use such an algorithm to compute exact results for some of our experimental comparisons.

For calculating the exact significance of a table under the hypergeometric model, Pagano and Taylor-Halvorsen [PTH81] have proposed a shortcut that typically avoids generating every table, but in the worst case may still be exhaustive. They make use of the structure of the hypergeometric distribution, and their trick does not apply under the uniform model. (The hypergeometric or

Fisher-Yates model can be viewed as based on the following sampling scheme. Imagine a population of exactly N members, where for each j , $1 \leq j \leq n$ there are c_j elements of type j . For each i , $1 \leq i \leq m$, fill row r_i by sampling uniformly independently *without* replacement from among the population, and if an element of type j is picked, place it in column j of row i . This is different from the multinomial model in which we sample *with* replacement from a population in which a fraction c_j/N falls in category j .)

C.1.2 A Recursive Formula

Gail and Mantel [GM77] give a straightforward recursive formulation of M_{rc} . For example, for the $m \times 3$ case, they give

$$M(r_1, r_2, \dots, r_m; c_1, c_2) = \sum_{k_1, k_2} M(r_1, r_2, \dots, r_{m-1}; c_1 - k_1, c_2 - k_2),$$

where $M(r_1, r_2, \dots, r_m; c_1, c_2)$ denotes the number of tables with the prescribed row sums r_i , $1 \leq i \leq m$, and the prescribed column sums c_j , $1 \leq j \leq n$; the third column sum is implicit, since it is determined by $c_3 = N - c_1 + c_2$. The sum on the right runs over the values k_1 and k_2 with

$$0 \leq k_i \leq \min(r_m, c_i) \quad \text{and} \quad k_1 + k_2 \leq r_m.$$

While this might seem at first to suggest a quick way to compute M_{rc} in general, the generalization of this recurrence to arbitrary dimension and margins takes exponential time to compute.

C.1.3 Approximation Formulas

Diaconis and Efron [DE85] provide the following approximation for M_{rc} using a volume-times-density argument. They do not provide an explicit means of estimating the error. Let

$$w = \frac{1}{1 + mn/2N},$$

$$\bar{r}_i = \frac{1-w}{m} + \frac{wr_i}{N}, \quad \text{for } 1 \leq i \leq m$$

$$\bar{c}_j = \frac{1-w}{n} + \frac{wc_j}{N}, \quad \text{for } 1 \leq j \leq n$$

and

$$k = \frac{n+1}{n \sum_i \bar{r}_i^2} - \frac{1}{n}.$$

Then

$$M_{rc} \approx \left(\frac{2N + mn}{2} \right)^{(m-1)(n-1)} \left(\prod_{i=1}^m \bar{r}_i \right)^{n-1} \left(\prod_{j=1}^n \bar{c}_j \right)^{k-1} \frac{\Gamma(nk)}{\Gamma(n)^m \Gamma(k)^n}. \quad (\text{C.1})$$

They carefully motivate the approximation and give some empirical evidence that it is accurate. We use (C.1) in some comparisons in Chapter 5.

Good [Goo76] has suggested a very similar approximation, with slightly less 'tuning' to the parameters. Gail and Mantel [GM77] suggest a normal approximation.

Remark: The asymmetry between rows and columns in the approximation is an artifact of a procedural choice in their derivation, namely to use a certain distribution on the column margins conditioning on the row margins rather than the other way around. \square

C.1.4 Tables and Group Theory

The set Σ_{rc} and its cardinality M_{rc} arise in a variety of contexts connected with the symmetric group and its representations. We describe some examples below to show more reasons to be interested in this set, and also to suggest alternative algorithms to compute the size.

Double Cosets

Given a partition r of N , let Y_r be the subgroup of the symmetric group S_N consisting of all permutations that permute the first r_1 elements among themselves, the next r_2 elements among themselves, and so on; Y_r is called a **Young subgroup**. It is isomorphic to the direct product of the symmetric groups S_{r_1} . If Y_r and Y_c are two Young subgroups of S_N , one can define an equivalence relation \sim on S_N by

$$\pi \sim \sigma \text{ if and only if } \rho\pi\kappa = \sigma, \text{ some } \rho \in Y_r, \kappa \in Y_c.$$

The equivalence classes of this relation are called the **double cosets** of Y_r and Y_c in S_N . It is a classical fact that there is a 1-1 correspondence between the double cosets and the elements of Σ_{rc} . For a group-theoretic proof, the reader is referred to James and Kerber [JK81, Corollary 1.3.11]. The correspondence has the following algorithmic combinatorial interpretation. Consider N balls labelled 1 up to N . Color the first r_1 balls color 1, the next r_2 balls color 2, and so on. Let π be a permutation in S_N , viewed as an arrangement of the the balls by their numeric labels. Construct a table $T(\pi)$ as follows: Look at the first c_1 places in π and for each color i , count how many balls of color i occur in these places. Call these numbers $T(\pi)_{i,1}$, and write them in the corresponding entries along column 1 of the table $T(\pi)$. Then look at the next c_2 places in π and count how many balls of each color i occur, calling these numbers $T(\pi)_{i,2}$, writing them along column 2 of $T(\pi)$. Continuing in this way produces an $m \times n$ table, $T(\pi)$, whose i th row sum is r_i and j th column sum is c_j . Thus $T(\pi) \in \Sigma_{rc}$. It is not hard to check that every table in Σ_{rc} is $T(\pi)$ for some π , and that $T(\pi) = T(\sigma)$ if and only if π and σ are in the same double coset. Note that if we let π be the identity (or any fixed permutation), this algorithm also gives a way to construct an element of Σ_{rc} .

Induced Representations, Tensor Products

Let G be a finite group, $H \subset G$ a subgroup, $X = G/H$ the associated coset space, and $L(X)$ the vector space of all real-valued functions on X . The group G acts on $L(X)$ by left translation: $[sf](x) = f(s^{-1}x)$. This action gives the representation of G induced by the trivial representation of H (see [Ser77]), denoted $\text{Ind}_H^G(\text{triv})$. For $G = S_N$ with Young subgroup $H = Y_r$, these representations arise directly in the statistical analysis of “partially ranked data of shape r ” [Dia88] [Dia89].

The representations $\text{Ind}_{Y_r}^{S_N}(\text{triv})$ and $\text{Ind}_{Y_c}^{S_N}(\text{triv})$ decompose into irreducible representations. The dimension of the space of common irreducible components is called the intertwining number $I(r, c)$. Formally, this is the dimension of the space of linear maps from $\text{Ind}_{Y_r}^{S_N}(\text{triv})$ to $\text{Ind}_{Y_c}^{S_N}(\text{triv})$ which commute with the action of the group. A classical theorem of Mackey (see [JK81, p. 17]) states that

$$I(r, c) = M_{rc}.$$

Remark: Here is a related fact. If, instead, we consider $\text{Ind}_{Y_r}^{S_N}(\text{triv})$ and $\text{Ind}_{Y_c}^{S_N}(\text{alt})$, where ‘alt’ denotes the alternating representation, the resulting intertwining number is equal to the number of 0-1 matrices with row sums r and column sums c . If r is a partition of N , and $c = r^T$, the partition obtained by transposing the Ferrers diagram of r , the resulting intertwining number is 1. This unique common constituent, S^c , is an irreducible representation of S_N . As c varies over all partitions, each irreducible representation occurs as S^c once and only once. James and Kerber [JK81, pp. 34-36] use this as their basic construction of the irreducible representations of the symmetric group. \square

If $(\rho_1, V_1), (\rho_2, V_2)$ are linear representations of G , one can form the tensor product of the two representations. It is the set of matrices $\rho_1(s) \otimes \rho_2(s)$, for $s \in G$, and gives a representation of G on $V_1 \otimes V_2$. This is a basic construction of representation theory, and the need arises to understand how the tensor product decomposes.

If I^r denotes $\text{Ind}_{Y_r}^{S_N}(\text{triv})$, then

$$I^r \otimes I^c = \bigoplus_T I^s,$$

where the direct sum ranges over all tables in $s \in \Sigma_{rc}$, where the entries of the table are linearly ordered to form a partition of N . This is proved and discussed by James and Kerber [JK81, pp. 95-98]. The connection between tensor products and double cosets is a special case of a more general phenomenon discussed by Curtis and Reiner [CR81].

Symmetric Functions

A polynomial is called symmetric if it is invariant under every permutation of its variables. Let Λ_n be the ring of symmetric polynomials in n variables, x_1, x_2, \dots, x_n , with integer coefficients. This ring decomposes into a direct sum of subrings

$$\Lambda_n = \bigoplus_{k \geq 0} \Lambda_n^k,$$

where Λ_n^k consists of the homogeneous symmetric polynomials of degree k , together with the zero polynomial. The reader should refer to MacDonald [Mac79] for further information and background on material appearing throughout this section.

There are a number of well-known bases for the symmetric functions. The four most common ones are: the monomial symmetric functions m_λ , the elementary symmetric functions e_λ , and the complete symmetric functions h_λ , and the power sum symmetric functions p_λ . (Elements of these bases are conveniently indexed by partitions λ of n .)

A scalar product can be defined on Λ_n by requiring that the h and m bases should be dual, i.e.

$$\langle h_\lambda | m_\mu \rangle = \delta_{\lambda, \mu},$$

where δ is the Kronecker delta. Moreover, the power sum basis is orthogonal with respect to this inner product. We have

$$\langle p_\lambda | p_\mu \rangle = \delta_{\lambda, \mu} z_\lambda, \quad (C.2)$$

where if a_i is the multiplicity of i in the partition λ , then $z_\lambda = \prod_{i \geq 1} i^{a_i} a_i!$.

MacDonald [Mac79, Sec. I-6] shows that if r and c are partitions of N , then

$$\langle h_r | h_c \rangle = M_{rc}.$$

This suggests an algorithm for computing M_{rc} : Given r and c , express h_r and h_c in the power-sum basis. Multiply in this basis using the orthogonality relation C.2. The result is M_{rc} . It can be shown that this yields an $O(N^{\max(m, n)})$ algorithm.

Example C.1 (Magic Squares) Let $H_n(r)$ denote the number of $n \times n$ matrices each of whose row and column sums are (the same value) r . Such matrices are called **magic squares**, and were first analysed by MacMahon [Mac16]. This is a special case of determining M_{rc} . Stanley [Sta73] [Sta86, Prop. 4.6.19, p. 232] has shown that $H_n(r)$ is a polynomial in r of degree exactly $(n-1)^2$. This polynomial is known for $n \leq 6$. Being a polynomial of degree $(n-1)^2$, $H_n(r)$ is determined by its values on $(n-1)^2 + 1$ points. Stanley showed that

$$H_n(-1) = H_n(-2) = \cdots = H_n(-n+1) = 0,$$

and

$$H_n(-n-r) = (-1)^{n-1} H_n(r).$$

So, for example, if we calculate $H_n(i)$, for $1 \leq i \leq \binom{n-1}{2}$, we can determine the polynomial $H_n(r)$. The coefficients of $H_5(r)$ and $H_6(r)$ were found in this manner by Jackson and Van Rees [JR75] using symmetric function techniques. It may now be feasible to find $H_7(r)$. The high order coefficient of $H_n(r)$ is also interesting for another reason; it is the volume of the polytope of $n \times n$ doubly stochastic matrices. [Sta86]. \square

Kostka Numbers, RSSK Correspondence

There are other natural bases of the symmetric functions in which to try working. For example, the Schur functions s_λ are characterized by being orthonormal in the inner product defined in the previous section. This means

$$\langle s_\lambda | s_\mu \rangle = \delta_{\lambda, \mu}.$$

The coefficients giving the change of basis from the Schur functions to the complete symmetric functions h_λ are called **Kostka numbers**. In other words, the Kostka numbers $K_{\lambda, \mu}$ are the unique numbers satisfying the relation

$$h_\lambda = \sum_{\mu} K_{\lambda, \mu} s_\mu.$$

Macdonald [Mac79, p.57] gives a straightforward proof from these definitions that

$$M_{rc} = \sum_{\mu} K_{\mu, r} K_{\mu, c}.$$

D. E. Knuth extended ideas of Robinson, Schensted, and Schützenberger to obtain an elegant combinatorial interpretation of this identity in terms of an explicit constructive correspondence based on Young tableaux. For λ and μ partitions of n , a semi-standard Young tableau of shape λ and content μ is a Young tableau of shape λ containing μ_1 1's, μ_2 2's, and in general μ_i i 's. The contents must be arranged in the cells of the tableaux in nondecreasing order (left to right) along the rows and strictly decreasing order (down) each column. The Kostka number $K_{\lambda, \mu}$ is equal to the number of such tableaux. The RSSK correspondence gives a constructive mapping between tables with row sums r and column sums c and pairs of semi-standard Young tableaux of contents r and c . (See [Knu68, Vol. 3] and [Knu70].)

It may thus be possible to generate uniform random tables in Σ_{rc} by generating pairs of semi-standard Young tableaux in the uniform distribution. This still seems hard. A random walk on pairs of tableaux generated by some Bernoulli-Laplace-like process is conceivable, and the coupling techniques used in Chapter 2 might be of some utility.

C.1.5 Counting with Generating Functions

Let x_1, x_2, \dots, x_m and y_1, y_2, \dots, y_n be variables. Form the generating function

$$\prod_{i,j} (1 - x_i y_j)^{-1} = (1 + x_1 y_1 + (x_1 y_1)^2 + \dots)(1 + x_1 y_2 + (x_1 y_2)^2 + \dots) \dots (1 + x_m y_n + (x_m y_n)^2 + \dots). \quad (C.3)$$

One can verify that the coefficient of

$$x_1^{r_1} x_2^{r_2} \dots x_m^{r_m} y_1^{c_1} y_2^{c_2} \dots y_n^{c_n}$$

is exactly the number of $m \times n$ tables with row sums \mathbf{r} and column sums \mathbf{c} .

For example, the coefficient of $x_1^2 x_2 y_1 y_2 y_3$ is 3, which means that for $\mathbf{r} = (2, 1)$ and $\mathbf{c} = (1, 1, 1)$, the value of $M_{\mathbf{rc}}$ = 3. In fact, one can check that the three tables are

1	1	0	1	0	1	0	1	1
0	0	1	0	1	0	1	0	0

The coefficients in the generating function can be expressed as contour integrals involving the generating function and this leads to some asymptotic approximations. Good [Goo76] records some efforts in this direction.

Along more algorithmic lines, we can compute initial portions of such generating functions by multiplying polynomials. This does not give polynomial-time algorithms, but this method may be feasible for small m , n , and N . We outline this method here. We will use \mathbf{i} to denote (i_1, i_2, \dots, i_k) , and write a term $z_1^{i_1} z_2^{i_2} \dots z_k^{i_k}$ as $z^{\mathbf{i}}$. We may then write a polynomial as $f(z_1, z_2, \dots, z_k) = f(z) = \sum_{\mathbf{i}} a_{\mathbf{i}} z^{\mathbf{i}}$.

Two k variable polynomials of degree at most d in each variable can be multiplied using fast Fourier transforms on the group Z_{2d}^k in $O(k(2d)^k \lg d)$ arithmetic operations. To see this, suppose $f(z) = \sum_{\mathbf{i}} a_{\mathbf{i}} z^{\mathbf{i}}$ and $g(z) = \sum_{\mathbf{i}} b_{\mathbf{i}} z^{\mathbf{i}}$ are two polynomials. Then their product is $h(z) = \sum_{\mathbf{s}} c_{\mathbf{s}} z^{\mathbf{s}}$ where

$$c_{\mathbf{s}} = \sum_{\mathbf{i} + \mathbf{j} = \mathbf{s}} a_{\mathbf{i}} b_{\mathbf{j}},$$

is a convolution. This convolution can be computed by multiplying two Fourier transforms on the group Z_{2d}^k . Standard FFT algorithms on this group of order $(2d)^k$ can be used to compute the transforms, inverse-transform, and the pointwise multiplications in the stated $O((2d)^k \lg((2d)^k)) = O(k(2d)^k \lg d)$ arithmetic operations. (See for example [CLR90, Chap. 32].)

Let $d_* = \max(\{r_i \mid i \in [m]\} \cup \{c_j \mid j \in [n]\})$ be the maximum among all of the row and column sums r_i and c_j . Since degrees increase monotonically by addition when we multiply polynomials, and the term whose coefficient is $M_{\mathbf{rc}}$ in the generating function (C.3) has no variable of degree exceeding d_* , we can eliminate from consideration any term with a variable whose order exceeds d_* . Successively multiplying each of the mn polynomials, and discarding terms of excess degree, gives an initial segment of the generating function that is accurate in the coefficient we need. The value of $M_{\mathbf{rc}}$ can be computed in this way using $O(mn(m+n)(2d_*)^{m+n} \lg d_*)$ arithmetic operations on $O((2d_*)^k)$ numbers.

When this technique is feasible, it can also be applied to situations in which it is desired to count the elements T of $\Sigma_{\mathbf{rc}}$ satisfying additional linear constraints of the form

$$\sum_{i,j} a_{ij} T_{ij} = a.$$

The basic generating function for tables can then be augmented to

$$\prod (1 - s^{a_{ij}} z_i y_j)^{-1},$$

so that the coefficient of $s^a x^r y^c$ is the number of tables in Σ_{rc} satisfying the additional constraint. Additional variables can be used to add more constraints.

For example, the generating function for “ $n \times n$ magic squares with diagonal constraints” can be expressed with two additional variables s and t as

$$\prod_{i,j} (1 - s^{a_{ij}} t^{b_{ij}} x_i y_j)^{-1},$$

where $a_{ij} = 1$ precisely if $i = j$, $b_{ij} = 1$ precisely when $i + j = n$, and both are zero otherwise.

Such additional constraints arise naturally in various types of contingency table inference. For example, Agresti, Mehta, and Patel [AMP90] needed the number of tables with prescribed row and column sums and one additional global constraint where $a_{ij} = u_i v_j$ for prescribed values of u_i and v_j .

Remark: A different approach involving generating functions is discussed by Stanley [Sta86, Chapter 4], who suggests an approach to counting the integer lattice points within a convex polytope with rational vertices. Determining M_{rc} may be viewed as such a counting problem. (See our Sections 5.3 and 5.5.) In particular, Stanley uses his method to prove some results about the number of magic squares (See our Examples 5.21 and C.1). This approach may yield similar results for more general instances of Σ_{rc} . \square

C.2 Hardness of a Related Problem

We have not been able to determine the complexity of the counting problem for Σ_{rc} . However, in this section we show that a related problem is provably difficult.

Let Σ be a finite alphabet. If $z \in \Sigma^*$, we use $|z|$ to denote the length of z . If $L \subseteq \Sigma^*$, and w is a given word in Σ^* , define the extension language

$$(L|w) = \{z | wz \in L\}.$$

In other words $(L|w)$ is the set of extensions z of w , that result in a word wz in L .

Let $R \subseteq \Sigma^* \times \Sigma^*$ be a relation on words. We write $R(z, y)$ to denote $(z, y) \in R$, and $R(z)$ to denote the set $\{y | R(z, y)\}$. R is a p-relation if there are polynomials p and q such that

- the predicate $R(z, y)$ can be checked in deterministic time $p(|z|)$.
- if $y \in R(z)$ then $|y| \leq q(|z|)$.

For a relation $R(z, y)$, let $R(z) = \{y | R(z, y)\}$. We call $R(z)$ the solution set of R given z . The class NP can be defined as the class of sets that can be represented $\{z | R(z) \neq \emptyset\}$ for some p-relation R . If R is a p-relation, the question “Given z , is there a y such that $R(z, y)$?” is called the decision problem for the relation. (For further background see [GJ79].)

Similarly, one may pose the counting problem for a p -relation R by asking "Given x , how many y are there such that $R(x, y)$?" The class $\#P$ is the class of counting problems for p -relations, or, more formally, the class of nonnegative integer functions f over Σ such that $f(x) = |R(x)|$ for some p -relation R .

There are elements f in $\#P$ such that the existence of a polynomial time algorithm that computes f would imply the existence of polynomial time algorithms for each function in $\#P$. Such a counting problem f is called $\#P$ -complete; these problems are in a sense the hardest in $\#P$. We now show that the counting problem for a certain family of sets related to Σ_{rc} is $\#P$ -complete.

For a set of matrix coordinates $Z \subset [m] \times [n]$, the set Σ_{rc}^Z of contingency-tables with structural zeros at Z , is the subset of Σ_{rc} in which every table has only zeros at the entries (i, j) , for $(i, j) \in Z$.

Theorem C.2 *With r , c , Z as parameters, the counting problem for Σ_{rc}^Z is $\#P$ -complete. This holds even if the inputs are expressed in unary.*

Proof: First note that the counting function for the set Σ_{rc}^Z is clearly in $\#P$. To prove completeness, we give a reduction from the problem of computing the permanent. Given an $n \times n$ matrix A with 0-1 entries, the permanent of A is the number of perfect matchings in the bipartite graph with vertex sets $[n]$ and $[n]$ and adjacency matrix A . The problem of computing the permanent was shown to be $\#P$ -complete by Valiant [Val79].

Computing the permanent is just a special case of computing $|\Sigma_{rc}^Z|$. Suppose we are given the $n \times n$ adjacency matrix A for a bipartite graph G . Let $r = c = (1, 1, 1, \dots, 1)$ (of dimension n). Construct the set Z of pairs (i, j) for which $A_{ij} = 0$. Now it is easy to see that a table T is in Σ_{rc}^Z if and only if T is the adjacency matrix of a perfect matching in G . Thus $|\Sigma_{rc}^Z|$ is equal to the number of perfect matchings in G . The described reduction can clearly be done in log-space. This proves the theorem. \square

We conjecture that the counting problem for Σ_{rc} is also $\#P$ -complete, even with parameters in unary. That is, we think it is unlikely that there is an algorithm that exactly counts Σ_{rc} in time polynomial in m , n , and N .

Remark: The set Σ_{rc}^Z , like the set Σ_{rc} , arises naturally in the analysis of contingency tables where the row and column classifications that gave rise to the table preclude certain combinations. For example, suppose that from a certain study we construct a table that classifies subjects by sex along the rows, and by the incidence of various cancers along the columns. Then a table entry representing the number of males with uterine cancer would necessarily contain a zero. (Statisticians call such a zero structural.) \square

C.3 Approximate Counting using Sampling

The techniques of Chapter 5 can be applied to do approximate counting of Σ_{rc} . In an earlier remark (on page 104), we noted that the techniques of Dyer, Frieze, and Kannan [DFK89] can be applied to approximately count Σ_{rc} . In this section, we show how our own sampling methods from Chapter 5 provide a means to do approximate counting, although we provide no polynomial-time guarantees.

The problems of counting a set and uniformly sampling from that set are closely related. This relation has been made precise in a complexity theoretic sense by Jerrum, Valiant, and Vazirani in [JV86]. In order to keep this work essentially self-contained, we will present a brief synopsis of their main idea here. However, to get a complete understanding of what we say here, the reader should be familiar with their ideas already or study their paper concurrently.

To clarify the setting, first consider the following simplified situation. Suppose one has a set V and subset $H \subseteq V$, where $|H|$ is known, but $|V|$ is not. If we could get a good approximation to $p = |H|/|V|$, then we could approximate $|V| = |H|/p$. This p is the mean value of the indicator function of the subset H under the uniform distribution on V . So we could approximate p by sampling (we presented an efficient means of doing this in Chapter 3), and provided p is not too small, the number of samples required to get a reasonable estimate will not be excessive.

Similarly, suppose we can find in V a sequence of nested subsets $V = H_1 \supset H_2 \supset \dots \supset H_m$ where $|H_m| = 1$ or some other known value, where each ratio $p_i = |H_{i+1}|/|H_i|$ is not too small. Then using the fact that $|V| = \prod_{i=1}^{m-1} (1/p_i)$ we can approximate $|V|$ using approximations for the p_i . We will work with sets V that admit such a decomposition, where, in fact, each subset H_i in the nested sequence is just a 'smaller instance' of V .

We now define a class of sets that admit this kind of decomposition in a way that can be computed 'efficiently.' Let Σ be a finite alphabet. A relation $R \subseteq \Sigma^* \times \Sigma^*$ is (polynomially) self-reducible if

- There is a deterministic polynomial-time computable function $g : \Sigma^* \rightarrow \mathbb{N}$ such that if $R(x, y)$ then $|y| = g(x)$.
- There exists polynomial-time computable functions $\psi : (\Sigma^*)^2 \rightarrow \Sigma^*$ and $\sigma : \Sigma^* \rightarrow \mathbb{N}$ such that for all $x, w \in \Sigma^*$.

$$\begin{aligned} \sigma(x) &\leq c \lg |x|, \text{ for some constant } c, \\ g(x) > 0 &\text{ implies } \sigma(x) > 0, \\ \text{and } |\psi(x, w)| &\leq |x|, \end{aligned}$$

and such that for all $x \in \Sigma^*$, if $y = wz$ with $|y| = g(x)$ and $|w| = \sigma(x)$ then

$R(x, wz)$ if and only if $R(\psi(x, w), z)$.

This tells us that $R(z, wz)$ can be determined by first computing $\psi(z, w)$, and then determining $R(\psi(z, w), z)$. We can also write this last condition as:

$$(R(z)|w) = R(\psi(z, w)).$$

That is, each extension language $(R(z)|w)$ can be expressed in terms of the original relation R on some smaller instance $\psi(z, w)$. Since $|w| = \sigma(z) = O(\ln |z|)$ the entire solution set $R(z)$ can be expressed as the disjoint union of a polynomial number of solution sets of the same relation on smaller instances. Many interesting sets can be expressed as the solution sets of self-reducible p-relations. (For the definition of a p-relation, see Section C.2.) We give one example below. Many others exist, including, for example, the perfect matchings in a bipartite graph, satisfying assignments of CNF or DNF formulae, and independent sets of given size in a graph.

Example C.3 (Spanning Trees) The relation $R(z, y) = "y \text{ is a spanning tree of the graph } z"$ is a self-reducible p-relation. Let z represent a graph with n nodes, (we will consider z to be an adjacency list, where each node index is a $\ell = \lceil \lg n \rceil$ bit integer). We will represent spanning trees as lists of $n - 1$ edges, each edge being a pair of node indices. So every solution y has length $|y| = g(z) = 2\ell(n - 1)$. Let $\sigma(z) = 2\ell$, so that the first $\sigma(z)$ characters of a solution y will represent the first edge in the list. For a string y of length $g(z)$, write $y = wz$, where $|w| = \sigma(z)$ and w represents one edge. Let $\psi(z, w)$ represent the result of contracting the edge w in z , (merging the vertices at the two ends of w and erasing any resulting multiple edges). This yields a smaller graph $|\psi(z, w)| \leq 2\ell(n - 2) \leq |z|$. Now note that $R(z, wz)$ if and only if $R(\psi(z, w), z)$. That is, $y = wz$ is a spanning tree of z if and only if, when we 'contract' the edge w in z , z represents a spanning tree of the resulting $\psi(z, w)$. We can also write $(R(z)|w) = R(\psi(z, w))$. \square

Using essentially the same reasoning as we outlined in the early paragraphs of this section, Jerrum, Valiant, and Vazirani show that if we can efficiently do near-uniform sampling in these sets then we can efficiently count these sets to within small relative error. In particular they give the following theorem.

Theorem C.4 (Counting using Sampling) (Adapted from [JV86]) *Let R be a self-reducible p-relation that is also in P . Suppose that we have an algorithm that*

- *takes input $z \in \Sigma^*$, and ϵ , $0 < \epsilon \leq 1$,*
- *runs in time polynomial in $|z|$ and $\ln(1/\epsilon)$,*
- *and generates a random sample $y \in R(z)$ whose distribution is within ratio $1 + \epsilon$ of the uniform distribution.*

Then, given an z , we can compute a count C such that C approximates $|R(z)|$ within ratio $1 + \epsilon$ with probability at least $1 - \delta$. Moreover, this computation can be done in time polynomial in $|z|$, $1/\epsilon$ and $\ln(1/\delta)$.

Remark: The requirement that the relation R is in P can be dropped if a slightly different notion of near-uniform sampling is used or if we are only concerned with inputs x for which $R(x)$ is nonempty. See [JV86]. \square

We will now show that we can cast Σ_{rc} in a self-reducible form. For any table $F \in \Sigma_{\text{rc}}$ and ordered pair $(k, l) \in [m] \times [n]$, let $[\Sigma_{\text{rc}}|F; (k, l)]$ denote the set of all tables T in Σ_{rc} with

$$T_{ij} = F_{ij} \text{ whenever } i < k \text{ and whenever } i = k \text{ and } j < l.$$

In other words, $[\Sigma_{\text{rc}}|F; (k, l)]$ is the subset of Σ_{rc} tables whose entries match the table F in all positions strictly preceding (k, l) in the lexicographic order. (We will use the symbols $>$ (and \succeq) to denote this order relation.) Notice that we have the following properties for any $F \in \Sigma_{\text{rc}}$: (a) $F \in [\Sigma_{\text{rc}}|F; (k, l)]$, (b) $[\Sigma_{\text{rc}}|F; (1, 1)] = \Sigma_{\text{rc}}$, and (c) for any $(k, l) > (m - 1, n - 1)$ we have $[\Sigma_{\text{rc}}|F; (k, l)] = \{F\}$, since all remaining entries are then determined by the sum constraints.

If we use $F_{k \leftarrow i}$ to denote the table obtained from F by setting $F_{ki} = i$, then we may write

$$[\Sigma_{\text{rc}}|F; (k, l)] = \bigcup_{0 \leq i \leq N} [\Sigma_{\text{rc}}|F_{k \leftarrow i}; \text{succ}(k, l)],$$

where $\text{succ}(k, l)$ is the ordered pair that is the immediate successor of (k, l) in the lexicographic order. Note that some of the sets on the right may be empty. This relation expresses the decomposition needed to show that $[\Sigma_{\text{rc}}|F; (k, l)]$ is polynomially self-reducible in the parameters m , n , and N .

Now we will show that we can use essentially the same random walk to draw samples from $[\Sigma_{\text{rc}}|F; (k, l)]$ as we used to draw samples from Σ_{rc} .

Algorithm C.5 (Random walk on $[\Sigma_{\text{rc}}|F; (k, l)]$) For a given k and l , modify Algorithm 5.5 so that in steps 2 and 3, it chooses only amongst values of i_1, i_2, j_1 and j_2 such that $(i_1, j_1) \succeq (k, l)$. That is replace those steps with the following.

2', 3' Uniformly choose a pair of rows i_1 and i_2 , $i_1 < i_2 \leq m$, and a pair of columns j_1 and j_2 , $j_1 < j_2 \leq n$, such that $(i_1, j_1) \succeq (k, l)$.

Theorem C.6 The random walk generated by the Algorithm C.5, and started on any $F \in \Sigma_{\text{rc}}$ is ergodic and has uniform stationary distribution on $[\Sigma_{\text{rc}}|F; (k, l)]$.

Proof: The walk always remains within $(\Sigma_{\text{rc}}|F, k, l)$ since no entries in positions preceding (k, l) are ever altered. That the walk is irreducible on $[\Sigma_{\text{rc}}|F; (k, l)]$ follows from the second clause of Lemma 5.3, which says, in effect, that there is a path between any two tables A and B in $[\Sigma_{\text{rc}}|F; (k, l)]$ does not involve coordinates lexicographically less than (k, l) . The arguments for symmetry and aperiodicity are the same as for the basic walk. \square

As before, we don't know that the walk converges in polynomial-time. However, under the supposition that it does, we have the next theorem

Theorem C.7 Suppose there exists a convergence guarantee for the random walk of Algorithm C.5 that is polynomial in m , n , N , and $1/\epsilon$. Then there is an algorithm that

1. takes inputs r , c , ϵ and δ
2. runs in time polynomial in m , n , N , $1/\epsilon$, and $\ln(1/\delta)$,
3. produces an approximate count C that is within ratio ϵ of M_{rc} with probability at least $1 - \delta$.

Proof: Since this set is polynomially self-reducible in the parameters m, n , and N , Theorem C.4 applies. Theorem C.6 combined with the hypothesized convergence guarantee provides the necessary means of near-uniform sampling from $[\Sigma_{rc}|F; (k, l)]$. To estimate M_{rc} , we estimate the cardinality of $\Sigma_{rc} = [\Sigma_{rc}|F; (1, 1)]$, where F is any table in Σ_{rc} . \blacksquare

Remark: Traditional methods of uniform sampling from a combinatorial set V rely on formulas for $|V|$, or implicitly suggest formulas for $|V|$. But in this setting $|V|$ is exactly what we don't know and wish to know. Broder [Bro86] made the valuable observation that sampling using a Markov chain avoids this problem. A number of other authors (e.g. [JS88, DFK89]) have since used Markov chains to do sampling for approximate counting. Note that the results of Chapter 3 on approximating mean values of indicator functions can be applied directly to improve the running times in all of these approximate counting algorithms as well. \square

Bibliography

- [AD86] D. Aldous and P. Diaconis. Strong uniform times and finite random walks. Technical Report 249, Department of Statistics, Stanford University, 1986.
- [AGM87] N. Alon, Z. Galil, and V. D. Milman. Better expanders and superconcentrators. *Journal of Algorithms*, 8:337–347, 1987.
- [AKL⁺79] R. Aleliunas, R. M. Karp, R. J. Lipton, L. Lovász, and C. Rackoff. Random walks, universal traversal sequences, and the complexity of maze problems. In *Proceedings of the 20th IEEE Symposium on the Foundations of Computer Science (FOCS)*, 1979.
- [AKS87] M. Ajtai, J. Komlós, and E. Szemerédi. Deterministic simulation of logspace. In *Proceedings of the 19th ACM Symposium on the Theory of Computing (STOC)*, 1987.
- [Ald87] D. Aldous. On the Markov chain simulation method for uniform combinatorial distributions and simulated annealing. *Probability in the Engineering and Informational Sciences*, 1:33–46, 1987.
- [Alo86] N. Alon. Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986.
- [AM85] N. Alon and V. D. Milman. λ_1 , isoperimetric inequalities for graphs, and superconcentrators. *Journal of Combinatorial Theory, Series B*, 38:73–88, 1985.
- [AMP90] A. Agresti, C. Mehta, and N. Patel. Exact inference for contingency tables with ordered categories. *Journal of the American Statistical Association*, 85:453–458, 1990.
- [Bab79] L. Babai. Spectra of Cayley graphs. *Journal of Combinatorial Theory, Series B*, 27:180–189, 1979.
- [Bab90] L. Babai. Local expansion of vertex-transitive graphs and random generation in finite groups. Technical report, Department of Computer Science, University of Chicago, 1990.
- [Bac87] E. Bach. Realistic analysis of some randomized algorithms. In *Proceedings of the 19th ACM Symposium on the Theory of Computing (STOC)*, 1987.

- [Ban80] C. Bande. *Isoperimetric Inequalities and Applications*. Pitman Advanced Publishing Program, 1980.
- [BD89] D. Bayer and P. Diaconis. Trailing the dovetail shuffle to its lair. Technical Report 329, Department of Statistics, Stanford University, 1989.
- [BFH75] Y. M. M. Bishop, S. E. Fienberg, and P. W. Holland. *Discrete Multivariate Analysis: Theory and Practice*. MIT Press, Cambridge Mass, 1975.
- [BGG90] M. Bellare, O. Goldreich, and S. Goldwasser. Randomness in interactive proofs. In *Proceedings of the 31st Annual IEEE Symposium on the Foundations of Computer Science (FOCS)*, pages 563–572, 1990.
- [Big74] N. L. Biggs. *Algebraic Graph Theory*. Cambridge Tracts in Mathematics, No. 67. Cambridge University Press, 1974. Unfortunately, this wonderful book is no longer in print.
- [BK88] A. Z. Broder and A. Karlin. Bounds on covering times. In *Proceedings of the 29th IEEE Symposium on the Foundations of Computer Science (FOCS)*, 1988.
- [BKL89] L Babai, W. M. Kantor, and A. Lubotsky. Small-diameter Cayley graphs for finite simple groups. *European Journal of Combinatorics*, 10:507–522, 1989.
- [BL84] G. Birkhoff and R. E. Lynch. *Numerical Solution of Elliptic Problems*. Number 6 in SIAM Studies in Applied Mathematics. SIAM, Philadelphia, 1984.
- [Bol86] B. Bollobás. *Combinatorics: Set Systems, Hypergraphs, Families of Vectors, and Combinatorial Probability*. Cambridge University Press, 1986.
- [BOP88] J. Baglivo, D. Olivier, and M. Pagano. Methods for the analysis of contingency tables with large and small cell counts. *Journal of the American Statistical Association*, 83(404):1006–1013, 1988.
- [Bro] A. Z. Broder. Couplings and strong-uniform times for the hypercube. Oral communication.
- [Bro86] A. Z. Broder. How hard is it to marry at random? (on the approximation of the permanent). In *Proceedings of the 18th ACM Symposium on the Theory of Computing*, 1986.
- [BW73] D. M. Boulton and C. S. Wallace. Occupancy of a rectangular array. *Computing*, 16(1):57–63, 1973.
- [Cal90] J. M. Calvin. *Stochastic Optimization Algorithms and Moment Formulas for Markov Chains*. PhD thesis, Stanford University, 1990.

[CDS80] D. M. Cvetković, M. Doob, and H. Sachs. *Spectra of Graphs: Theory and Applications*. Academic Press, New York, 1980.

[Che70] J. Cheeger. A lower bound for the smallest eigenvalue of the Laplacian. In *Problems in Analysis*. Princeton University Press, New Jersey, 1970.

[Chu60] K. L. Chung. *Markov Chains with Stationary Transition Probabilities*. Springer-Verlag, Berlin, 1960.

[CLR90] T. Cormen, C. Leiserson, and R. Rivest. *Introduction to Algorithms*. The MIT Electrical Engineering and Computer Science Series. McGraw-Hill/The MIT Press, 1990.

[CR81] C. W. Curtis and I. Reiner. *Methods of Representation Theory, with applications to finite groups and orders*. Wiley and Sons, New York, 1981. Two volumes.

[CW89] A. Cohen and A. Wigderson. Dispersers, deterministic amplification, and weak random sources. In *Proceedings of the 30th IEEE Symposium on the Foundations of Computer Science (FOCS)*, 1989.

[DE85] P. Diaconis and B. Efron. Testing for independence in a two-way table: New interpretations of the chi-square statistic. *The Annals of Statistics*, 13(3):845–874, 1985. Invited paper. Discusses the case for volume tests and gives asymptotic approximations for estimating volume-based significance of the Chi-square statistic; some opposing and supporting discussion ensues in pages following, and a rejoinder by the authors appears on page 905 of the same volume.

[DF88] P. Diaconis and J. A. Fill. Strong stationary times via a new form of duality. Technical Report 305, Department of Statistics, Stanford University, 1988.

[DFK89] M. Dyer, A. Frieze, and R. Kannan. A random polynomial time algorithm for approximating the volume of convex bodies. In *Proceedings of the 21st Annual ACM Symposium on the Theory of Computing (STOC)*, 1989.

[Dia88] P. Diaconis. *Group Representations in Probability and Statistics*, volume 11 of *Lecture Notes - Monograph Series*. Institute of Mathematical Statistics, 1988.

[Dia89] P. Diaconis. A generalization of spectral analysis with application to ranked data. *The Annals of Statistics*, 17(3):949–979, 1989. Content of the 1987 Wald Memorial Lectures.

[DS81] P. Diaconis and M. Shahshahani. Generating a random permutation with random transpositions. *Z. Wahrscheinlichkeitstheorie verw. Gebiete*, 57:159–179, 1981.

[DS87] P. Diaconis and M. Shahshahani. Time to reach stationarity in the Bernoulli-Laplace diffusion model. *SIAM Journal on Mathematical Analysis*, 18(1):208–218, January 1987.

- [DS89] P. Diaconis and D. Stroock. Geometric bounds for eigenvalues of Markov chains. Technical Report 325, Department of Statistics, Stanford University, 1989.
- [EE07] P. Ehrenfest and T. Ehrenfest. Über zwei bekannte Einwände gegen das Boltzmannsche H-Theorem. *Phys. Zeitschrift*, 8:311–314, 1907. Also see Feller, Volume 1, p. 121 and pp. 377ff.
- [ER61] P. Erdős and A. Rényi. On a classical problem of probability theory. *Magyar Tud. Akad. Matemat. Kutató Intézet. Közl.*, 6:215–219, 1961. Proves a limit law for the multiple coupon collecting problem when the desired number of complete sets of coupons is fixed.
- [Eve77] B. S. Everitt. *The Analysis of Contingency Tables*. Monographs on Applied Probability and Statistics. Chapman and Hall, London, 1977.
- [Fel70] W. Feller. *An Introduction to Probability Theory and its Applications*. Wiley and Sons, 1970. Two volumes: 3rd (Vol. 1) and 2nd (Vol. 2) revised editions. The first edition was printed in 1968.
- [Fie73] M. Fiedler. Algebraic connectivity of graphs. *Czechoslovakian Mathematics Journal*, 23:298–305, 1973.
- [Fil90] J. A. Fill. Eigenvalue bounds on convergence to stationarity for non-reversible Markov chains, with an application to the exclusion process. Technical report, The Johns Hopkins University, Department of Mathematical Sciences, 1990.
- [Fla82] L. Flatto. Limit theorems for some random variables associated with urn models. *The Annals of Probability*, 10(4):927–934, 1982. Asymptotic analysis of variables associated with multiple coupon collecting.
- [Flo] R. W. Floyd. Generating random samples without replacement, as sequences and sets. This unpublished note, dated April 1987, contains Floyd’s elegant random subset algorithm, and another for generating a random permutation of m elements out of $[n]$. Their only published appearance to date is in Jon Bentley’s column “Programming Pearls,” *Communications of the ACM*, 30(9):754–757, September 1987.
- [FOW85] L. Flatto, A. M. Odlyzko, and D.B. Wales. Random shuffles and group representations. *The Annals of Probability*, 13(1):154–178, 1985.
- [GC77] I. J. Good and J. F. Crook. The enumeration of arrays and a generalization related to contingency tables. *Discrete Mathematics*, 19:23–65, 1977.
- [GG81] O. Gabber and Z. Galil. Explicit constructions of linear-sized superconcentrators. *Journal of Computer and System Sciences*, 22:407–420, 1981.

- [GI84] P. W. Glynn and D. L. Iglehart. Recursive moment formulas for regenerative simulation. In J. Janssen, editor, *Proceedings of the International Symposium on Semi-Markov Processes and their Applications, Brussels*, 1984.
- [GI87] P. W. Glynn and D. L. Iglehart. A joint central limit theorem for the sample mean and regenerative variance estimator. *Annals of Operations Research*, 8, 1987.
- [GJ79] M. Garey and D. Johnson. *Computers and Intractability*. Freeman, 1979.
- [GL89] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, 1989. Second edition.
- [GLS88] M. Grötschel, L. Lovász, and A. Schriver. *Geometric Algorithms and Combinatorial Optimization*. Springer-Verlag, 1988.
- [GM77] M. Gail and N. Mantel. Counting the number of $r \times c$ contingency tables with fixed margins. *Journal of the American Statistical Association*, 72(360), 1977.
- [Goo76] I. J. Good. On the application of symmetric Dirichlet distributions and their mixtures to contingency tables. *Annals of Statistics*, 4:1159–1189, 1976.
- [Gri75] D. Griffeath. A maximal coupling for Markov chains. *Z. Wahrscheinlichkeitstheorie verw. Gebiete*, 31:95–106, 1975.
- [Gri78] D. Griffeath. Coupling methods for Markov processes. In *Advances in Mathematics Supplementary Studies: Studies in Probability and Ergodic Theory 2*, pages 1–43. Academic Press, 1978.
- [Hal82] Peter Hall. *Rates of Convergence in the Central Limit Theorem*. Number 62 in Research Notes in Mathematics. Pitman Publishing, Marshfield, Mass., 1982.
- [Han57] E.J. Hannan. The variance of the mean of a stationary process. *Royal Statistical Society Journal, Series B*, 19(2):282–5, 1957.
- [Han74] T. W. Hancock. Remark on algorithm 434. *Communications of the ACM*, 18:117–119, 1974.
- [Hil90] M. Hildebrand. *Rates of Convergence of Some Random Processes on Finite Groups*. PhD thesis, Department of Mathematics, Harvard University, 1990.
- [Hun74] T. W. Hungerford. *Algebra*. Number 73 in Graduate Texts in Mathematics. Springer-Verlag, New York, 1974.

- [Igl78] D. L. Iglehart. The regenerative method for simulation analysis. In Chandy K. M. and R. T. Yeh, editors, *Current Trends in Programming Methodology*, volume III: Software Engineering. Prentice-Hall, New Jersey, 1978.
- [ILL89] R. Impagliazzo, L. Levin, and M. Luby. Pseudo-random generation from one-way functions. In *Proceedings of the 21st ACM Symposium on the Theory of Computing (STOC)*, pages 12–24, 1989.
- [IZ89] R. Impagliazzo and D. Zuckerman. How to recycle random bits. In *Proceedings of the 30th IEEE Symposium on the Foundations of Computer Science (FOCS)*, 1989.
- [JK77] N. L. Johnson and S. Kotz. *Urn Models and their Application*. Wiley, New York, 1977.
- [JK81] G. D. James and A. Kerber. *The Representation Theory of the Symmetric Group*, volume 16 of *Encyclopedia of Mathematics and its Applications*. Addison-Wesley, Reading, Mass., 1981.
- [JM85] S. Jimbo and A. Maruoka. Expanders obtained from affine transformations. In *Proceedings of the 17th ACM Symposium on the Theory of Computing (STOC)*, pages 88–97, 1985.
- [JR75] D. M. Jackson and G. H. J. Van Rees. The enumeration of generalized double stochastic non-negative integer square matrices. *SIAM Journal on Computing*, 4:475–477, 1975. Gives the polynomials $H_5(r)$ and $H_6(r)$ for counting 5×5 and 6×6 ‘magic squares’ with sums r .
- [JS88] M. Jerrum and A. Sinclair. Approximating the permanent. Technical Report CSR-275-86, University of Edinburgh, Dept. of Computer Science, 1988.
- [JS90] M. Jerrum and A. Sinclair. Polynomial-time approximation algorithms for the Ising model. Technical Report CSR-1-90, University of Edinburgh, Dept. of Computer Science, 1990.
- [JV86] M. Jerrum, L. Valiant, and V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science*, 43:169–188, 1986.
- [Kar68] S. Karlin. *A First Course in Stochastic Processes*. Academic Press, New York, 1968. Second printing.
- [Knu68] D. E. Knuth. *The Art of Computer Programming*. Addison-Wesley, 1968. Three volumes. Volume 1 first appeared 1968. The material on Young tableaux appears in Volume 3 (1973).

- [Knu70] D. E. Knuth. Permutations, matrices, and generalized Young tableaux. *Pacific Journal of Mathematics*, 34:709–727, 1970.
- [KR88] H. Karloff and P. Raghavan. Randomized algorithms and pseudo-random numbers. In *Proceedings of the 20th ACM Symposium on the Theory of Computing (STOC)*, 1988.
- [LPS86] A. Lubotsky, R. Phillips, and P. Sarnak. Explicit expanders and the Ramanujan conjecture. In *Proceedings of the 18th ACM Symposium on the Theory of Computing (STOC)*, pages 240–245, 1986.
- [LS90] L. Lovász and M. Simonovits. The mixing rate of Markov chains, an isoperimetric inequality, and computing the volume. Technical Report Preprint 27, The Mathematical Institute of the Hungarian Academy of Sciences, 1990.
- [Mac16] P. A. MacMahon. *Combinatory Analysis*. Cambridge University Press, 1916. Reprinted in 1960 by Chelsea, New York as one volume.
- [Mac79] I. G. MacDonald. *Symmetric Functions and Hall Polynomials*. Clarendon Press, Oxford, 1979.
- [Mar72] D. L. March. Exact probabilities for $r \times c$ contingency tables. *Communications of the ACM*, 15:991–992, 1972.
- [Mat85] P. Matthews. *Covering Problems for Random Walks on Spheres and Finite Groups*. PhD thesis, Department of Statistics, Stanford University, 1985. Available as Technical Report No. 234.
- [MGB74] A. Mood, F. Graybill, and D. Boes. *Introduction to the Theory of Statistics*. McGraw Hill, 1974. Third Edition.
- [MM64] M. Marcus and H. Minc. *A Survey of Matrix Theory and Matrix Inequalities*. Allyn and Bacon, Boston, 1964.
- [MO63] L. E. Moses and R. V. Oakland. *Tables of Random Permutations*. Stanford University Press, 1963.
- [MPS88] C. R. Mehta, N. R. Patel, and P. Senchaudhuri. Importance sampling for estimating exact probabilities in permutational inference. *Journal of the American Statistical Association*, 83(404):999–1005, 1988.
- [MRR⁺53] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equations of state calculation by fast computing machines. *Journal of Chemical Physics*, 21:1087–1091, 1953.

- [MV88] M. Mihail and V. Vazirani. On the magnification of 0-1 polytopes. Unpublished manuscript., 1988.
- [NS60] D. J. Newman and L. Shepp. The double Dixie cup problem. *American Mathematical Monthly*, 67(1):58–61, 1960. Analyzes the expected number of coupons required in multiple coupon collecting, when a constant number of complete sets is desired.
- [PTH81] M. Pagano and K. Taylor-Halvorsen. An algorithm for finding the exact significance levels of $r \times c$ contingency tables. *Journal of the American Statistical Association*, 76(376):931–4, 1981.
- [PW60] L.E. Payne and H. F. Weinberger. An optimal Poincaré inequality for convex domains. *Arch. Rational Mech. Anal.*, 5:286–292, 1960.
- [Ser77] J.-P. Serre. *Linear Representations of Finite Groups*. Springer-Verlag, 1977. Translation of the French edition of 1971.
- [SJ87] A. Sinclair and M. Jerrum. Approximate counting, uniform generation, and rapidly mixing Markov chains. Technical Report CSR-241-87, University of Edinburgh, Dept. of Computer Science, 1987.
- [Sta73] R. P. Stanley. Linear homogeneous Diophantine equations and magic labelings of graphs. *Duke Math Journal*, 40:607–632, 1973.
- [Sta86] R. P. Stanley. *Enumerative Combinatorics*. Wadsworth and Brooks/Cole, Monterey, California., 1986. Only the first volume has appeared.
- [Tho86] H. Thorisson. On maximal and distributional coupling. *Annals of Probability*, 14:874–876, 1986.
- [Val79] L. G. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8:189–201, 1979.

**END
FILMED**

DATE:

4-17-96

NTIS